Redbooks
ibm.com/redbooks

IBM TS4500 R8 Tape Library Guide

Ole Asmussen

Robert Beiderbeck

Albrecht Friess

Hans-Günther Hörhammer

Khanh Ngo

Jesus Eduardo Cervantes Rolon

Fabian Corona Villarreal

Larry Coyne

Storage

IBM

Redbooks

International Technical Support Organization

IBM TS4500 R8 Tape Library Guide

March 2022

SG24-8235-09

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 the IBM TS4500 R8 Tape Library.

© Copyright International Business Machines Corporation 2014, 2021, 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

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

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.

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

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.

AIX®
FICON®
Guardium®
IBM®
IBM Research®
IBM Security™
IBM Spectrum®
IBM Z®
IBM z Systems®
NPS®
OS/400®
ProtecTIER®
RACF®
Redbooks®
Redbooks (logo) ®
System z®
Tivoli®
XIV®
z Systems®
z/OS®

The following terms are trademarks of other companies:

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.

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

Preface

The IBM® TS4500 (TS4500) tape library is a next-generation tape solution that offers higher storage density and better integrated management than previous solutions.

This IBM Redbooks® publication gives you a close-up view of the new IBM TS4500 tape library. In the TS4500, IBM delivers the density that today's and tomorrow's data growth requires. It has the cost-effectiveness and the manageability to grow with business data needs, while you preserve investments in IBM tape library products.

Now, you can achieve a low per-terabyte cost and high density, with up to 13 PB of data (up to 39 PB compressed) in a single 10 square-foot library by using LTO Ultrium 9 cartridges or 11 PB with 3592 cartridges. The TS4500 offers the following benefits:

- Support of the IBM Linear Tape-Open (LTO) Ultrium 9 tape drive: Store up to 1.04 EB 2.5:1 compressed per library with IBM LTO 9 cartridges.
- High availability: Dual active accessors with integrated service bays reduce inactive service space by 40%. The Elastic Capacity option can be used to eliminate inactive service space.
- Flexibility to grow: The TS4500 library can grow from the right side and the left side of the first L frame because models can be placed in any active position.
- Increased capacity: The TS4500 can grow from a single L frame up to another 17 expansion frames with a capacity of over 23,000 cartridges.
- High-density (HD) generation 1 frames from the TS3500 library can be redeployed in a TS4500.
- Capacity on demand (CoD): CoD is supported through entry-level, intermediate, and base-capacity configurations.
- Advanced Library Management System (ALMS): ALMS supports dynamic storage management, which enables users to create and change logical libraries and configure any drive for any logical library.
- Support for IBM TS1160 while also supporting TS1155, TS1150, and TS1140 tape drive.

  The TS1160 gives organizations an easy way to deliver fast access to data, improve security, and provide long-term retention, all at a lower cost than disk solutions. The TS1160 offers high-performance, flexible data storage with support for data encryption.

  Also, this enhanced fifth-generation drive can help protect investments in tape automation by offering compatibility with existing automation. Store up to 1.05 EB 3:1 compressed per library with IBM 3592 cartridges.
- Integrated TS7700 back-end Fibre Channel (FC) switches are available.
- Up to four library-managed encryption (LME) key paths per logical library are available.

This book describes the TS4500 components, feature codes, specifications, supported tape drives, encryption, new integrated management console (IMC), command-line interface (CLI), and REST over SCSI (RoS) to obtain status information about library components.

You learn how to accomplish the following tasks:

- Improve storage density with increased expansion frame capacity up to 2.4 times, and support 33% more tape drives per frame.
- Manage storage by using the ALMS feature.

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

- Improve business continuity and disaster recovery with dual active accessor, automatic control path failover, and data path failover.
- Help ensure security and regulatory compliance with tape-drive encryption and Write Once Read Many (WORM) media.
- Support IBM LTO Ultrium 9, 8, 7, 6, and 5, and IBM TS1160, TS1155, TS1150, and TS1140 tape drives.
- Provide a flexible upgrade path for users who want to expand their tape storage as their needs grow.
- Reduce the storage footprint and simplify cabling with 5U or 10U of rack space on top of the library.
- Optimize connectivity with FC, Ethernet, and SAS interface attachments.
- Simplify user access to data stored on tape by way of IBM Spectrum® Archive.
- Cyber resilient technology with physical air gap.

This publication is for anyone who wants to understand more about the IBM TS4500 tape library. It is particularly suitable for IBM clients, IBM Business Partners, IBM specialist sales representatives, and technical specialists.

Authors

This book was produced by an international team working with the IBM Tucson Redbooks center.

Ole Asmussen is a Product Field Engineer for IBM Storage based in Hamburg, Germany. He has supported IBM Tape for over 20 years, working in various support level and supporting all IBM Tape products (3494, 3584, TS7700, and ProtecTier). He joined IBM as a Customer Service Representative in 1999 and was responsible for different accounts in Germany. In 2001, he joined the Support Center in Mainz, Germany.

Robert Beiderbeck is an IBM Senior Accredited Specialist for the Product Service profession, and works as a Team Leader for the EMEA RMSS Product Field Engineering (PFE) group in Germany. He joined IBM in 1989 as a Customer Service Representative for Mainframe (IBM Z Systems®) and worked later in the Z Systems Hardware Support Center, responsible for Server, ESCON, and IBM FICON® Switches and Infrastructure. Robert then became a Storage Specialist for Tape Libraries, VTS, and Drives in the field before he joined second-level support in 2002. His areas of expertise include Enterprise Tape Systems, such as TS7700, TS3500, 3494, VTS (P2P), and 3590/3592 Drives and Control Units.

Albrecht Friess is an IBM Remote Technical Support (RTS) Specialist and works as a Subject Matter Expert (SME) for IBM EMEA Tape Storage, leveraging the close collaboration with product development, support centers, and other business units to provide seamless problem resolution and drive continuous quality improvements. After completing his studies with a degree in engineering, he joined IBM in 1996. After different roles, he joined second-level EMEA Tape Storage Hardware Support in 2012. His area of expertise includes Enterprise Tape Systems, such as TS3500 and TS4500, and attached LTO and Jaguar tape drive devices. In addition to his work in the remote support center, he supports customers across Europe on-site with complex installations or problems. Albrecht also is frequently asked to hold training courses for customers and customer service representatives. He further acts as a focal point for certain premium storage customers.

Hans-Günther Hörhammer is pre-sales support and services specialist. He joined IBM in 1999. Hans supported and delivered services for various products and solutions like DS8K, NAS Gatwéway, IBM XIV®. He developed HA/DR solutions for Oracle and SAP. Presales support for Tape libraries in Open System environment is his latest area of expertise.

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 Enterprise Edition (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 integrating IBM Spectrum Archive EE into customers' production environments.

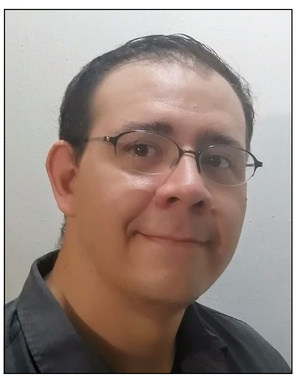

Jesus Eduardo Cervantes Rolon is an IBM Software Test Engineer in Guadalajara, Mexico. He holds a bachelor's degree in Electronics and Computer Engineering from CETI (Mexico) and he is a certified ISTQB FL Tester. He has been with IBM since 2012. During the last 7 years, he has been collaborating with the IBM next generation of storage management systems, the IBM TS4500 tape library in the Product Test area (focused in Functional Verification Test) where he has served as the team leader for the last 3 years.

Fabian Corona Villarreal is a Test Engineer for the IBM Mexico Software Lab, he joined IBM in 2004 and he has been being involved in multiple test projects since then. He started working on Storage products in 2010 when he joined the IBM ProtecTIER® team. There, he served as a test engineer, project leader and team leader during the different releases of the product. In 2014, he moved to the physical tapes area as the RAS team lead for the 3584, 3576, 3573, 3572, and 3555 tape libraries. Fabian holds a bachelor's degree in Computer Science from the Universidad de Guadalajara. He also is a certified ISTQB Advanced Level Test Manager.

Larry Coyne is a Redbooks Project Leader in the 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 client relationship management, quality assurance, development management, and support management for IBM storage management software.

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

Carlos Sandoval Castro
Larry Hogan
Lee Jesionowski
Javier Justo
Jeremy Tudgay
Rick West
Roberta Winston
IBM Systems

Special thanks to Michael Engelbrecht, Simon Browne, and Illarion Borisevich for their contributions to the prior editions.

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 2 - 6 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:

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:

  ibm.com/redbooks

- Send your comments by email:

  redbooks@us.ibm.com

- Mail your comments:

  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

Summary of changes for *IBM TS4500 R8 Tape Library Guide*, SG24-8235, as created or updated in March 2022.

This section describes the technical changes that were 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.

December 2021, Tenth Edition

This update reflects the addition of new and changed information.

New information

TS4500 Release 8 added the following functions:

- Added support for Support of the IBM Linear Tape-Open (LTO) Ultrium 9 tape drive (2.6, “IBM LTO Ultrium 9 tape drive (Model 3588 F9C, F9S, and S9C)” on page 141)
- Added “Open Recommended Access Order” on page 144
- Updated Chapter 3, “Encryption” on page 173 with IBM Security™ Guardium® Key Lifecycle Manager (formerly IBM Security Key Lifecycle Manager)
- Added/updated sections in Chapter 4, “TS4500 management graphical user interface” on page 185
  - 4.6.4, “SAS Ports” on page 274
  - 4.11, “IBM Net Promoter Score Feedback” on page 323
  - “Remote authentication”
- Updated Chapter 5, “Command-line interface” on page 325
- Updated Chapter 6, “TS4500 REST API” on page 391 and added the following sections:
  - “Introduction of Work Items” on page 393
  - “Moving diagnostic cartridges to I/O station” on page 406
  - “Reset node card” on page 467
  - “Slots” on page 474
  - “Cartridge Movement Request - Move To Drive” on page 483
  - “Cartridge Movement Request - Move To IO Station” on page 484
  - “Cartridge Movement Request - Move Cartridge To Slot” on page 485

Changed information

- Updated library firmware
- Updated tables, figures, examples, and reference links

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

July 2021, Ninth Edition (Update)

This update reflects the addition of new and changed information.

Changed information

- Updated IBM Spectrum Archive icon in 2.11, “IBM Spectrum Archive” on page 162.
- Updated 2.2, “IBM TS1160 tape drive” on page 105 and 2.3, “IBM TS1155 and TS1150 tape drive” on page 114 with Streaming Lossless Data Compression (SLDC).

October 2020, Ninth Edition (Update)

This update reflects the addition of new and changed information.

New information

TS4500 Release 7 added the following functions:

- Added support for the 3592 model 60S tape drive that provides a dual-port 12 Gb SAS (Serial Attached SCSI) interface for host attachment (Chapter 1, “IBM TS4500 tape library” on page 1, Chapter 2, “TS4500 Ultrium Linear Tape-Open and 3592 tape drives” on page 77, and Chapter 4, “TS4500 management graphical user interface” on page 185)

Changed information

Corrected cleaning cartridge label naming convention in , “Cleaning cartridges” on page 136.

June 2020, Ninth Edition

This update reflects the addition of new and changed information.

New information

TS4500 Release 7 added the following functions:

- Support for rear door open detection and reporting (each frame requires plant feature code 4892)
- Activity log on GUI System Summary
- REST support for data cartridges, frames, and reports
- Added an option for a new 5U top rack frame, 3584 Model TR2 providing an extra 5U of rack space on any frame in a library without requiring more floor space.

Changed information

Reorganization and consolidation of frame cartridge slot capacity.

January 2020, Eighth Edition

This update reflects the addition of new and changed information.

New information

The book includes the following new information:

- TS4500 REST API - REST over SCSI (RoS) commands
- New CLI commands for Service actions

Changed information

Updated Table 1-14 on page 47 (January 22, 2020).

December 2018, Seventh Edition

This update reflects the addition of new and changed information.

New information

The book includes the following new information:

- TS1150 model 60E
- TS1160 model 60F
- New media JE, JM, and JV

This update reflects the addition of new and changed information.

May 2018, minor update

Clarification in section “Mainframe-ready” on page 26 and “Feature codes for the frame models of the TS4500 tape library” on page 57 that the TS4500 must have two integrated 16 Gb switches installed for IBM z/OS® systems with an attached TS7700 support.

January 2018, Sixth Edition

This update reflects the addition of new and changed information.

New information

The book includes the following new information:

- LTO Ultrium 8 tape drive (Model 3588 F8C).
- `LTO 8 Type M cartridge (M8)`: The LTO Program is introducing a new capability with LTO-8 drives. The ability of the LTO-8 drive to write 9 TB on a brand new LTO-7 cartridge instead of 6 TB as specified by the LTO-7 format. Such a cartridge is called an LTO-7 initialized LTO-8 Type M cartridge.

This update reflects the addition of new and changed information.

August 2017, Fifth Edition

New information

The book includes the following new information:

- TS1155 model 55E
- TS1155 model 55F
- New commands for Ethernet drives

July 2016, Fourth Edition

This update reflects the addition of new and changed information.

New information

The book includes the following new information:

- High availability with dual active accessor and Elastic Capacity option
- Mainframe-ready for TS7700 attachment
- Integrated TS7700 Backend Switches
- Support for external TSSC/IMC
- Flexible growth options with new flex track design
- Encryption options

November 2015, Third Edition

This update reflects the addition of new and changed information.

New information

The book includes the following new information:

- IBM TS1070 tape drive (Model 3588 F7C)
- Up to 128 tape drives
- Scalability to 18 frames
- Redeployment of S24 and S54 frames from TS3500 to TS4500
- Mixed media types within the same TS4500 library
- Automatic media verification
- Flexible remote authentication
- Primary control system failover
- SNMP query configuration

January 2015, Second Edition

This update reflects the addition of new and changed information.

New information

The book includes the following new information:

- IBM TS1155 and TS1150 tape drive
- IBM TS1100 tape drives for the TS4500 tape library
- IBM LTO Ultrium tape drives for the TS4500 tape library

1

IBM TS4500 tape library

The IBM TS4500 tape library is a storage solution that is designed to help midsize and large enterprises respond to storage challenges. Several of these challenges include the high data volume, growth in data centers, increasing cost of data center storage footprints, difficulty migrating data across vendor platforms, and increased complexity of IT training and management as staff resources shrink.

The TS4500 tape library combines reliable, automated tape handling and storage with high performance in an open systems and enterprise environment. Incorporating the IBM Linear Tape-Open (LTO) Ultrium tape and 3592 drives, the TS4500 tape library offers outstanding retrieval performance with typical cartridge move times of less than 3 seconds.

The TS4500, installed with the High Availability (HA) option, provides dual active accessors for redundancy and can double the robot performance during tape move operations. The HA option on TS4500 has no dedicated service bays and provides the Elastic Capacity option for the use of the storage slots in these integrated service bays.

The TS4500 tape library can be deployed as a single frame library and upgraded to a maximum of 18 frames, with a combination of LTO and 3592 frames. This single frame library can be partitioned into multiple logical libraries. This feature makes the TS4500 tape library an excellent choice for consolidating tape workloads from multiple heterogeneous open systems servers.

The TS4500 protects investment by providing for redeployment of S24 and S54 frames from the TS3500 onto the TS4500.

The library provides outstanding reliability and redundancy through the provision of redundant power supplies in each drive frame, control and data path failover, dual grippers within the cartridge accessor, and dual active accessors. Library and drive firmware can be upgraded nondisruptively, without interrupting normal operations.

Encryption is available for all supported drives. The following encryption methods are supported:

- Application-managed encryption (AME)
- Library-managed encryption (LME)
- System-managed encryption (SME) for IBM z/OS TS7700 support

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

TS4500 Release 8 added the following functions:

- 3588-F9C (LTO9 FC Multi Mode) Drive Support
- 3588-F9S (LTO9 FC Single Mode) Drive Support
- 3588-S9C (LTO9 SAS) Drive Support
- TLS support for GKLM
- REST support for accessor speed adjustment
- REST support for moving cartridges (`moveToSlot`, `moveToDrive`, and `moveToIOStation`)
- REST support to request specific log files from the library
- REST support for reporting frame door open/close transitions as attributes rather than states
- REST support for querying the state of all slots or using a location parameter
- REST support for `scannerFailed` state for Accessors resource
- REST support for `accessible` attribute for Drive and Cartridges resources
- REST support for `internalAddress` attribute for Cartridges resources and ability to move all cartridges types to I/O station
- REST support for new tasks of `calibrateFrame` and `calibrateLibrary`
- REST support for new attributes of GET /v1/reports/drives for drive clean

TS4500 Release 7 added the following functions:

- TS1160 model 60S
- Support for rear door open detection and reporting (each frame requires plant feature code 4892)
- Activity log on GUI System Summary
- REST support for data cartridges, frames, and reports
- An option for a new 5U top rack frame, 3584 Model TR2, which provides an extra 5U of rack space on any frame in a library without requiring more floor space.

TS4500 Release 6 added the in-band method for sending REST API commands and receiving HTTP responses by using SCSI Write Buffer and Read Buffer commands, respectively. The method is called REST over SCSI (RoS).

TS4500 Release 5 added the following functions:

- TS1160 model 60E
- TS1160 model 60F
- Fibre Channel connectivity report

TS4500 Release 4 added the following functions:

- TS1155 model 55E
- TS1155 model 55F

TS4500 Release 4.1 added the LTO Ultrium 8 tape drive (Model 3588 F8C) function.

TS4500 Release 3 added the following functions, which are described in this chapter:

- HA with dual active accessor and Elastic Capacity option
- Mainframe-ready for TS7700 attachment
- Integrated TS7700 Backend Switches
- Support for external TSSC/Integrated management console (IMC)

- Flexible growth options with new flex track design
- Support for 4 EKM servers on each logical library

With TS4500 Release 2, the following new functions are supported over and above the first release:

- Automatic media verification
- Flexible remote authentication
- Primary control system failover
- Mixed media types within the same TS4500 library
- Scalability to 18 frames
- Up to 128 tape drives
- SNMP query configuration
- Redeployment of S24 and S54 frames from TS3500 to TS4500

This chapter includes the following topics:

- 1.1, “Overview of IBM TS4500 tape library” on page 4
- 1.2, “TS4500 product description” on page 5
- 1.3, “TS4500 tape library components” on page 31
- 1.4, “Feature codes for the TS4500” on page 57
- 1.5, “Host platforms and device drivers” on page 71
- 1.6, “Specifications” on page 72

1.1 Overview of IBM TS4500 tape library

The IBM TS4500 is a highly scalable, stand-alone tape library that provides high-density tape storage and high-performance, automated tape handling for open systems, and enterprise environments.

Figure 1-1 shows a seven-frame version of the TS4500 tape library. An individual library can consist of 1 L frame and up to 17 expansion frames. It also can include up to 128 tape drives with more than 23,000 tape cartridges, as shown in Figure 1-3.

Figure 1-1 TS4500 base with six expansion frames

The TS4500 tape library provides the following capabilities:

- High availability dual active accessors with integrated service bays to reduce inactive service space by 40%. The Elastic Capacity option can be used to completely eliminate inactive service space.
- All of the frames include high-density (HD) slot technology.
- Extra HD2 frame models can be placed in any active position so that the library can grow from the right and left side of the first L frame.
- HD generation 1 frames from the TS3500 library can be redeployed into a TS4500. These frames must be installed to the right of the Lx5 frame. Feature Code (FC) 1742 must be installed on each frame before they can exist in a TS4500 library string.
- In dual accessor configurations, the integrated service bays reduce the number of unused storage columns in a dual accessor library from 22 to 14. While an accessor is in service, the media columns in that area are not available to the second accessor.
- Only Dx5 frames without the I/O station and Sx5 frames are supported as frame one in a dual accessor configuration. The rightmost frame can be an Lx5, Dx5, or Sx5 model.
- Advanced Single Deep Cell technology.
- Integrated management console (IMC) with support for an external IBM TotalStorage System Console (TSSC) and IMC.
- Web-based user interface for improved usability.
- Updated control system.
- Input/output (I/O) magazine to allow individual cartridge handling to be performed independently of the library.

- Top-rack space to house extra tape solution components within the library footprint.
- Support for HD2-compatible models of the following tape drives:
  - TS1160 (3592 60E, 3592 60F, and 3592 60S)
  - TS1155 (3592 55E and 3592 55F)
  - TS1150 (3592 EH8)
  - TS1140 (3592 EH7)
  - LTO Ultrium 9 (3588 F9C, 3588 F9S and 3588 S9C)
  - LTO Ultrium 8 (3588 F8C)
  - LTO Ultrium 7 (3588 F7C)
  - LTO Ultrium 6 (3588 F6C)
  - LTO Ultrium 5 (3588 F5C)
- Integrated TS7700 back-end Fibre Channel switches.
- Up to four library-managed encryption (LME) key paths per logical library.
- The TS4500 tape library is available with the following tape drives, frame models, and feature options to meet your specific needs:
  - Advanced Library Management System (ALMS)
  - Ability to attach multiple simultaneous heterogeneous servers
  - Remote management with the TS4500 management GUI or the TS4500 command-line interface (CLI)
  - Remote monitoring by using Simple Network Management Protocol (SNMP), email, or syslog
  - SNMP query configuration
  - Media health verification
  - Multipath architecture
  - Drive and media exception reporting
  - Host-based path failover
  - Up to 288 I/O slots (36 I/O slots standard for LTO libraries and 32 I/O slots standard for 3592 libraries with extra I/O slots that are available as a feature add-on for all D25 and D55 frames)

1.2 TS4500 product description

The IBM TS4500 tape library (Machine Type 3584) is a modular tape library that consists of a high-density base frame and up to 17 high-density expansion frames. The frames join side by side and can grow to the left or right of the base frame. All frames are supported by up to two cartridge accessors. You can install a single-frame base library (see Figure 1-2 on page 6) and grow it to 18 frames, which tailors the library to match your system capacity requirements.

The supported combinations of frames, tape drives, and their capabilities are listed in Table 1-1.

Table 1-1 TS4500 tape library capabilities

Models	Drives in frames	Maximum cartridges	Maximum native capacity
L25, D25, S25, and S24	3592 tape drives	17,550	351 petabytes (PB)
L55, D55, S55, and S54	LTO tape drives	23,170	417 petabytes (PB)
Note: The maximum native capacity figures are based on library configurations of one base frame with all LTO 9 or TS1160 tape drives, and 17 storage-only HD frames.			

The base TS4500 tape library is shown in Figure 1-2.

Figure 1-2 Base TS4500 tape library

TS4500 expansion frames can be wrapped with custom images, as shown in Figure 1-3.

Figure 1-3 TS4500s with image wrapped expansion frames

Eight types of frames are supported in the current TS4500 tape library range. Each frame is identified by a three-character model number (L25, D25, L55, D55, S25, S55, S24, and S54), which describes the nature of the frame.

The TS4500 tape library is built from a single frame model that is called the *base frame*. The scalability of the library allows an increase in capacity by adding up to 17 frames, which are called *expansion frames*. The frames join side by side and can grow to the left or right of the base frame. All frames can be supported by a dual cartridge single accessor, or by dual active accessors, with the HA feature installed. The TS4500 tape library can contain a mix of 3592 and LTO frames.

The TS4500 tape library supports first generation (S54 and S24) frames (HD1) and second-generation high-density (HD2) frames. HD2 frames, as with the first-generation HD1 frames, offer increased capacity without increasing the frame size or required floor space, by using high-density storage slots for tape cartridges.

In addition, HD2 frames provide the following enhancements:

- HD2 frames can be installed in the leftmost position of the library (frame number 1).
- Drive-capable HD2 frames support up to 16 HD2-compatible tape drives (3588 F9C, F9S, S9C, F8C, F7C, F6C, F5C, 3592 EH7, EH8, 55E, 55F, 60E, 60F, and 60S) when positioned as frame number 2 or higher.

Generation 1 HD frames from the TS3500 (Model S24 and S54) can be redeployed into a TS4500. These HD1 frames must be installed to the right of the Lx5 frame and require FC 1742 to be ordered for each S24 or S54 (Sx4) frames before they can exist in a TS4500 library string.

The L25 and L55 (Lx5) frames and D25 and D55 (Dx5) frames are HD2, drive-capable frames, which means that they contain high-density cartridge storage slots, and slots to house up to 16 tape drives. The S25, S55 (Sx5) HD2 frames and the S54, S24 (Sx4) HD1 frames are storage-only frames, which means that they contain high-density cartridge storage slots, but no tape drives. All HD frames provide internal light-emitting diode (LED) lighting.

The TS4500 also supports adding a top rack frame. The top rack, 3584 Model TR1, provides an extra 10U of rack space and 3584 Model TR2 provides an extra 5U of rack space on any frame in a library without requiring more floor space.

The frames that are supported by the library and their specific media type and capacity are listed in Table 1-2.

Table 1-2 TS4500 tape library frame models

Frame model	Type	Media type	Capacity		Other
			Frame position 1	Frame position 2+	
L25	Base frame	3592	Up to 12 tape drives and 552 storage slots	Up to 16 tape drives and 660 storage slots	► Equipped with two I/O stations and two 16-slot magazines ► Optionally equipped with top rack (Model TR1 or TR2)
L55	Base frame	LTO	Up to 12 tape drives and 732 storage slots	Up to 16 tape drives and 882 storage slots	► Equipped with two I/O stations and two 18-slot magazines ► Optionally equipped with top rack (Model TR1 or TR2)
D25	Expansion frame	3592	Up to 12 tape drives and 592 storage slots	Up to 16 tape drives and 740 storage slots	► Optionally quipped with two I/O stations and two 16-slot magazines ► Optionally equipped with top rack (Model TR1 or TR2)
D55	Expansion frame	LTO	Up to 12 tape drives and 776 storage slots	Up to 16 tape drives and 970 storage slots	► Optionally equipped with two I/O stations and two 18-slot magazines ► Optionally equipped with top rack (Model TR1 or TR2)
S25	Storage-only expansion frame	3592	800 storage slots	1,000 storage slots	Optionally equipped with top rack (Model TR1 or TR2)
S55	Storage-only expansion frame	LTO	1,056 storage slots	1,320 storage slots	Optionally equipped with top rack (Model TR1 or TR2)
S24	Storage-only expansion frame	3592	Not supported	1,000 storage slots	Optionally equipped with top rack (Model TR1 or TR2)
S54	Storage-only expansion frame	LTO	Not supported	1,320 storage slots	Optionally equipped with top rack (Model TR1 or TR2)

Capacity-on-Demand

In the TS4500 tape library, the physical capacity (or total storage slots) is composed of licensed and unlicensed capacity. When the number of assigned cartridges reaches the licensed capacity, more cartridges cannot be assigned to a logical library until a cartridge is removed, a CoD feature is purchased, or frames are added to the library.

The Intermediate, Base, and High-Density Capacity-on-Demand features provide license keys so that you can enable more storage slots in the frames of the TS4500 tape library.

Intermediate and base capacity on demand

Use the Intermediate and Base Capacity on Demand (CoD) features to increase the initial (entry) capacity of the base frames (models L25 and L55) of the TS4500 tape library.

The initial (entry) capacity of the L25 and L55 frames is 100 storage slots. You can purchase CoD features to increase the amount of available licensed capacity.

The Intermediate CoD feature (FC 1643) adds 100 slots, which increases the usable capacity of the L25 and L55 frames to 200 slots. The Base CoD feature (FC 1644) adds 200 slots, which increases the usable capacity of the L25 and L55 frames to 400 slots. FC 1644 is referred to as Full CoD with the TS4500 tape library.

High-Density capacity on demand

Use the High Density (HD) Capacity on Demand (CoD) license key to enable the full high-density capacity of the Lx5, Dx5, and Sx5 frames in the TS4500 tape library. Enabling the CoD on the TS4500 is a non-disruptive process.

The initial (entry) capacity of the Lx5 frames is 100 slots. The Intermediate and Base CoD features can increase the usable capacity up to 400 slots. The HD CoD features add 150 to more than 450 slots, depending on frame position and configuration. The Base CoD feature (FC 1644) is a prerequisite for installing an HD CoD feature on an Lx5 frame.

The initial (entry) capacity of the Dx5 frames is 500 slots. The initial (entry) capacity of the S25 frame is 600 slots and the S55 frame is 660 slots. The HD CoD features can add 50 - 660 slots, depending on frame position and configuration.

The potential capacity by frame model is listed in Table 1-3.

Table 1-3 Potential capacity by frame model

Frame Model	Licensed feature	F1 slots available	F2+ slots available
L25	Entry	100	100
	Intermediate	200	200
	Base	400	400
	HD CoD	552	660
L55	Entry	100	100
	Intermediate	200	200
	Base	400	400
	HD CoD	732	882
D25	Base	500	500
	HD CoD	592	740
D25 with IO	Base	500	500
	HD CoD	552	660
D55	Base	500	500
	HD CoD	776	970
D55 with IO	Base	500	500
	HD CoD	732	882

Frame Model	Licensed feature	F1 slots available	F2+ slots available
S25	Base	600	600
	HD CoD	800	1000
S55	Base	660	660
	HD CoD	1056	1320
S24	Base	NA	600
	HD CoD	NA	1000
S54	Base	NA	660
	HD CoD	NA	1320

1.2.1 TS4500 tape library frames for IBM LTO Ultrium Fibre Channel drives

The TS4500 tape library models L55 and D55 integrate the HD2 versions of the LTO-9, LTO-8, LTO-7, LTO-6, and LTO-5 tape drives. The TS4500 models S55 and S54 are high capacity storage-only frames for LTO cartridge slots.

The Model L55 frame includes the frame control assembly with two power supplies (for redundancy), an optimized dual-gripper cartridge accessor, on-demand storage slot capacity, and two I/O stations with two 18-slot magazines.

TS4500 tape library Model L55

The L55 frame can be installed on its own as a complete library enclosure (as shown in Figure 1-4) or up to 17 expansion frames can attach to it. This frame provides the major library components for the entire library, whether it has a single frame or multiple frames. It also provides cartridge storage capacity for LTO media, and can be equipped with the HD2 versions of the LTO-9, 8, 7, 6, and 5 dual-ported drives that facilitate 8 Gbps Fibre Channel connectivity.

HD2 expansion frames can be added to the left or right of the L55 frame. HD1 frames can be added only to the right side of L55 frame (see Figure 1-4).

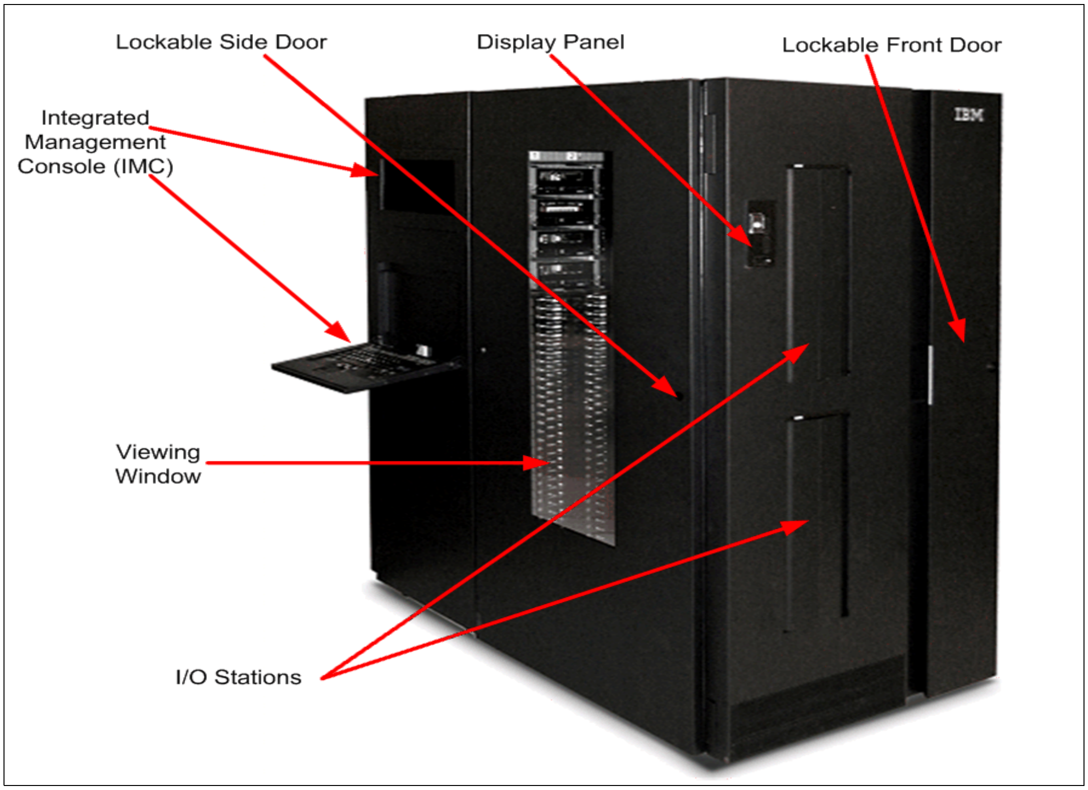

Figure 1-4 TS4500 tape library L55/L25 base frame

The number of LTO cartridge storage slots ranges 100 - 882. With the minimum configuration, 100 slots are available for use. The maximum of 882 slots is physically installed and accessed by adding CoD license keys.

Many CoD feature codes exist for the L55 frame. The number of available slots depends on the frame position.

The Intermediate Capacity feature (FC 1643) gives a maximum total number of usable cartridge slots of 200. This feature is a prerequisite for the Base Capacity on Demand (FC 1644), which gives the maximum capacity of 400 cartridge slots. FC 1644 is required to attach an optional expansion frame. Both FC 1643 and FC 1644 are prerequisites to install the HD CoD for L55 (FC 1648), which gives the maximum capacity of 730 - 882 slots.

Depending on the frame position, a maximum of 16 LTO drives can be installed. Five generations of HD2-compatible LTO drives exist: the LTO Ultrium 9 tape drive (Model 3588 F9C, F9S, S9C), the LTO Ultrium 8 tape drive (Model 3588 F8C), LTO Ultrium 7 tape drive (Model 3588 F7C), LTO Ultrium 6 tape drive (Model 3588 F6C), and the LTO Ultrium 5 tape drive (Model 3588 F5C), which can be installed in the L55 frame. Drive slots are fixed. Adding drives to the L55 frame does not affect the number of available storage slots.

Figure 1-5 shows the drive slots and HD slots.

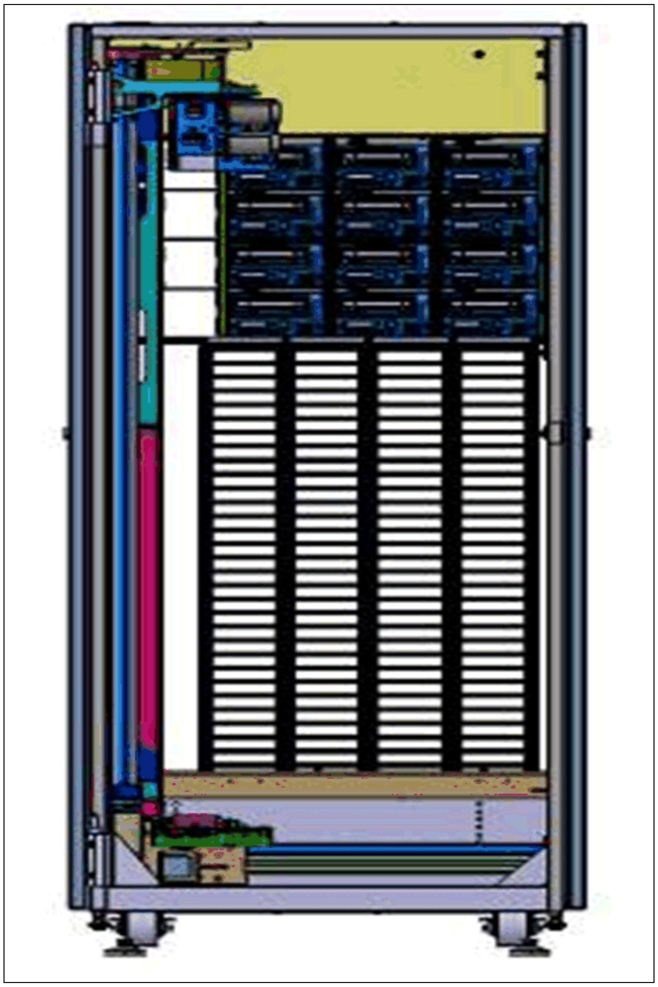

Figure 1-5 L55/L25 frame internal

When CoD features are installed, the position and configuration of the frame affect the total available capacity of the L55. The available storage capacity, which is based on the frame positions and configurations and capacity for each Tier, is listed in Table 1-4.

Table 1-4 Quantity of storage slots in the L55 frame

Licensed feature	F1 slots available	F2+ slots available
Entry	100	100
Intermediate	200	200
Base	400	400
HD CoD	732	882

The L55 frame included two I/O stations as standard. Each I/O station houses a cartridge magazine that allows individual cartridge handling to be performed independently of the tape library. The cartridge magazine for each I/O station on LTO frames can hold up to 18 cartridges, which provides a total of 36 I/O slots.

The TS4500 tape library Model L55 imports or exports cartridges from the library, without requiring reinventory or interruption of library operations. The lockable library door can be opened for bulk-loading LTO tape cartridges. Reinventory of the cartridges in tier 0 and tier 1 is performed in less than 60 seconds per frame, each time that the library door is closed. A bar code reader that is mounted on the autochanger is used to scan the cartridge bar code labels during inventory.

Important: If a bulk load is performed, the top two rows on tier 1 (drive side wall) must remain empty to allow for the initial inventory. Place only the cartridges in the frame that has the front door open. Do not insert cartridges into slots in an adjacent frame.

On an HD frame after initial inventory, the inventory checks tier 1 bar code labels only. It checks the other tier labels only if tier 1 changed.

A door lock is included to restrict physical access to the cartridges in the library. A door open sensor is equipped to prevent accessor movement while the door is open.

Note: The left and right side doors also contain the door lock and open sensor.

Included in the L55 frame is the IMC, which is a built-in platform for tools that are used to manage the TS4500 tape library. The IMC, which includes an LCD panel and a keyboard with a touchpad or track point, can be mounted on either end of your TS4500 tape library.

For more information about the IMC and other components, see 4.1, “Integrated management console” on page 186.

TS4500 tape library Model D55

The D55 frame, as shown in Figure 1-6, features the same footprint as the Model L55.

Figure 1-6 TS4500 model D55/L55

The D55 frame cannot be installed on its own. It must be connected to a library with a base frame. A maximum of 18 frames, including the L55 frame, can be connected, as shown in Figure 1-3 on page 7.

> **Note:** The combined number of D55 or D25 drive frames that can be installed in a TS4500 library is limited to seven.

The number of extra LTO cartridge storage slots per D55 frame ranges is 500 - 970. With the minimum configuration, only 500 slots are available for use. More slots can be enabled by installing a capacity on demand (CoD) license key.

The base capacity on a D55 frame gives the maximum capacity of 500 cartridge slots. FC 1644 must be installed on the L55 frame with FC 9002 or FC 9003, and the corresponding prerequisite feature code, to attach a D55 expansion frame, as described in 1.4, “Feature codes for the TS4500” on page 57. The HD CoD for D55 (FC 1650) gives the maximum capacity of 730 - 970 slots, depending on the frame position.

Depending on the frame position, the maximum number of LTO drives that can be installed is 16. As with the L55 frame, the following generations of HD2-compatible LTO drives can be installed in the D55 frame:

- The LTO Ultrium 9 tape drive (Model 3588 F9C, F9S and S9C)
- The LTO Ultrium 8 tape drive (Model 3588 F8C)
- The LTO Ultrium 7 tape drive (Model 3588 F7C)

- The LTO Ultrium 6 tape drive (Model 3588 F6C)
- The LTO Ultrium 5 tape drive (Model 3588 F5C)

Drive slots are fixed. Adding drives to the D55 frame does not affect the number of available storage slots.

Figure 1-6 on page 14 shows the D55 frame with drive and HD slots.

Two extra I/O stations can be installed in any Dx5 expansion frame by ordering FC 1652. This feature installs two I/O stations in a drive expansion frame. Each extra pair of I/O stations increases the maximum insert/eject throughput for the library. The maximum cartridge capacity for expansion frames with two I/O stations is reduced by 88 cartridges for the Model D55.

The position and configuration of the frame, the number of I/O slots, and the installation of capacity on demand (CoD) features all affect the total available storage capacity of the D55.

The available storage capacity, which is based on possible frame positions and configurations and capacity for each Tier, is listed in Table 1-5.

Table 1-5 Quantity of storage slots in the D55 frame

Special frame considerations	Licensed feature	F1 slots available	F2+ slots available
D55 with no IOs	Base	500	500
D55 with no IOs	HD CoD	776	970
D55 with IOs	Base	500	500
D55 with IOs	HD CoD	732	882

1.2.2 TS4500 tape library frames for IBM 3592 drives

The TS4500 tape library models L25 and D25 integrate the TS1160, TS1155, TS1150, and TS1140 tape drives. The TS4500 Model S25 is a high-capacity, storage-only frame for 3592 slots.

The Model L25 frame includes the Frame Control Assembly (FCA) with two power supplies (for redundancy), an optimized dual-gripper cartridge accessor, on-demand storage slot capacity, and two I/O stations with two 16-slot magazines.

TS4500 tape library Model L25

The L25 can be installed on its own as a complete library enclosure (as shown in Figure 1-4 on page 11) or it can have up to 17 expansion frames that are attached to it. This frame provides the major library components for the whole library, whether it has single or multiple frames. It also provides cartridge storage capacity for 3592 media and can be equipped with TS1160 (3592 model 60F) facilitating dual-ported 16 Gbps Fibre Channel connectivity, or TS1155, TS1150, and TS1140 (3592 models 55F, EH8, and EH7), which facilitates dual-ported 8 Gbps Fibre Channel connectivity.

The TS1160 (3592 model 60E) provides dual 10 or 25 Gb Ethernet host attachment interface. The TS1155 (3592 model 55E) provides dual 10 Gb Ethernet host attachment interface, which is optimized for cloud-based and large, open-compute environments.

The TS1160 (3592 model 60S) provides a dual-port 12 Gb SAS (Serial Attached SCSI) interface for host attachment. This drive brings more versatility to businesses with substantial storage, backup, and archiving demands with a cost-competitive communications interface to help simplify storage management and system performance.

The HD2 expansion frame can be added to the left or right of the L25 frame. HD1 frames can be added only to the right side of the L25 frame.

The Intermediate Capacity feature (FC 1643) gives a maximum of 200 usable cartridge slots. This feature is a prerequisite for the Base Capacity on Demand feature (FC 1644), which gives the maximum capacity of 400 cartridge slots. FC 1644 is required to attach an optional expansion frame. FC 1644 is a prerequisite to install the HD CoD for L25 (FC 1647), which offers the maximum capacity of 550 - 660 slots.

Depending on the frame positions, a maximum of 16 3592 drives can be installed. Four generations of HD2-compatible 3592 drives, the TS1160 (3592 60E, 60F, and 60S), TS1155 (3592 55F and 55E), TS1150 (3592 E08), and TS1140 (3592 E07) tape drives, are supported in the L25 frame. Drive slots are fixed. Adding drives to the L25 frame does not affect the number of available storage slots.

Figure 1-7 shows an L25 frame drive and HD slots.

Figure 1-7 TS4500 model D25/L25

The position and configuration of the frame and the installation of CoD features affect the total available capacity of the L25. The available storage capacity, which is based on possible frame positions and configurations and capacity for each Tier, is listed in Table 1-6 on page 17.

Table 1-6 Quantity of storage slots in the L25 frame

Licensed feature	F1 slots available	F2+ slots available
Entry	100	100
Intermediate	200	200
Base	400	400
HD CoD	552	660

The L25 frame includes two I/O stations as standard. Each I/O station houses a cartridge magazine that allows importing or exporting cartridges from the library without requiring reinventory or an interruption of library operations. The cartridge magazine for each I/O station on 3592 frames can hold up to 16 cartridges, which provides a total of 32 I/O slots.

The lockable library door can be opened for bulk-loading IBM LTO Ultrium tape cartridges. Reinventory of the cartridges in tier 0 and tier 1 is performed in less than 60 seconds per frame each time that the library door is closed. A bar code reader that is mounted on the gripper is used to scan the cartridge bar code labels during inventory.

Important: If a bulk load is performed, the top two rows on tier 1 (drive side wall) must remain empty to allow for the initial inventory.

On an HD frame, the inventory checks tier 0 and tier 1 bar code labels only, and the inventory checks the other tier labels only if tier 1 changed.

A door lock is included to restrict physical access to cartridges in the library. A door open sensor also is equipped to prevent accessor movement while the door is open.

Included in the L25 frame is the IMC, which is a built-in platform for tools that are used to manage the TS4500 tape library. The IMC, which includes an LCD panel and a keyboard with a touchpad, can be mounted on either end of your TS4500 tape library. For more information about the IMC and other components, see Chapter 4, “TS4500 management graphical user interface” on page 185.

TS4500 tape library Model D25

The D25 frame, as shown in Figure 1-6 on page 14, features the same footprint as the Model L25. The D25 frame cannot be installed on its own. It must be connected to a library with a base frame. A maximum of 18 frames, including the L25 frame, can be connected, as shown in Figure 1-3 on page 7.

Important: The combined number of allowed D55 or D25 frames in a TS4500 library is limited to seven.

The number of extra 3592 cartridge storage slots per D25 frame is 500 - 740. With the minimum configuration, only 500 slots are available for use.

More slots can be enabled by installing a CoD license key.

The base capacity on a D25 frame gives the maximum capacity of 500 cartridge slots.
FC 1644 must be installed on the L55 frame with FC 9002 or FC 9003 and the corresponding prerequisite FC to attach a D25 expansion frame, as described in 1.4, “Feature codes for the TS4500” on page 57.

The HD CoD for D25 (FC 1649) gives the maximum capacity of 660 - 740 slots, depending on the frame position.

Depending on the frame position, a maximum of 16 3592 drives that can be installed. The following drives are supported in the D25 frame:

- Four generations of HD2-compatible 3592 drives
- TS1160 (3592 60E, 60F, and 60S)
- TS1155 (3592 55E and 55F)
- TS1150 (3592 E08)
- TS1140 (3592 E07)

Drive slots are fixed. Adding drives to the D25 frame does *not* affect the number of available storage slots.

Figure 1-7 on page 16 shows the drive slots and HD slots.

Two extra I/O stations can be installed in any Dx5 expansion frame by ordering FC 1652. This feature installs two I/O stations in one expansion frame. Each extra pair of I/O stations increases the maximum insert/eject throughput for the library. The maximum cartridge capacity for expansion frames with two I/O stations is reduced by 80 cartridges for the Model D25 frame.

The position and configuration of the frame, the number of I/O slots, and the installation of CoD features all affect the total available storage capacity of the D55. The available storage capacity, based on possible frame positions and configurations and capacity for each Tier, is listed in Table 1-7.

Table 1-7 Quantity of storage slots in the D25 frame

Special frame considerations	Licensed feature	F1 slots available	F2+ slots available
D25 with no IOs	Base	500	500
D25 with no IOs	HD CoD	592	740
D25 with IOs	Base	500	500
D25 with IOs	HD CoD	552	660

1.2.3 TS4500 tape library storage-only HD frames

In this section, we describe the storage-only HD frames that are offered by the TS4500 tape library.

Model S55 and S25

The IBM TS4500 includes the Model S25 frame and the Model S55 frame, which are high density (HD) version 2 storage-only expansion frames, as shown in Figure 1-8.

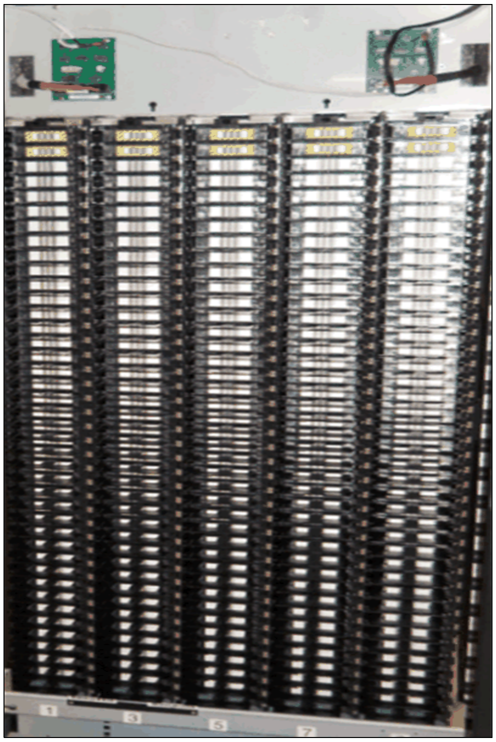

Figure 1-8 Sx5/Sx4 frame

These frames are designed to increase storage capacity greatly without increasing the frame size or required floor space.

The HD slots contain tape cartridges in a tiered architecture. The cartridge, which is immediately accessible in the HD slot, is a tier 1 cartridge (behind that tier is tier 2, and so on).

The maximum tier in an LTO HD slot is tier 5. The maximum tier in a 3592 HD slot is tier 4 because the 3592 tape cartridge is slightly longer than the LTO cartridge. The single-deep slots on the door side of HD frames are referred to as *tier 0 slots*.

A side view of the inside of an HD frame is shown on the left side of Figure 1-9. A top-down view of one row of an HD frame with cartridges in tier 0 (door side), 1 (Drive side), 2, 3, 4, and 5 is shown on the right side of Figure 1-9. Tier 5 is for LTO frames only.

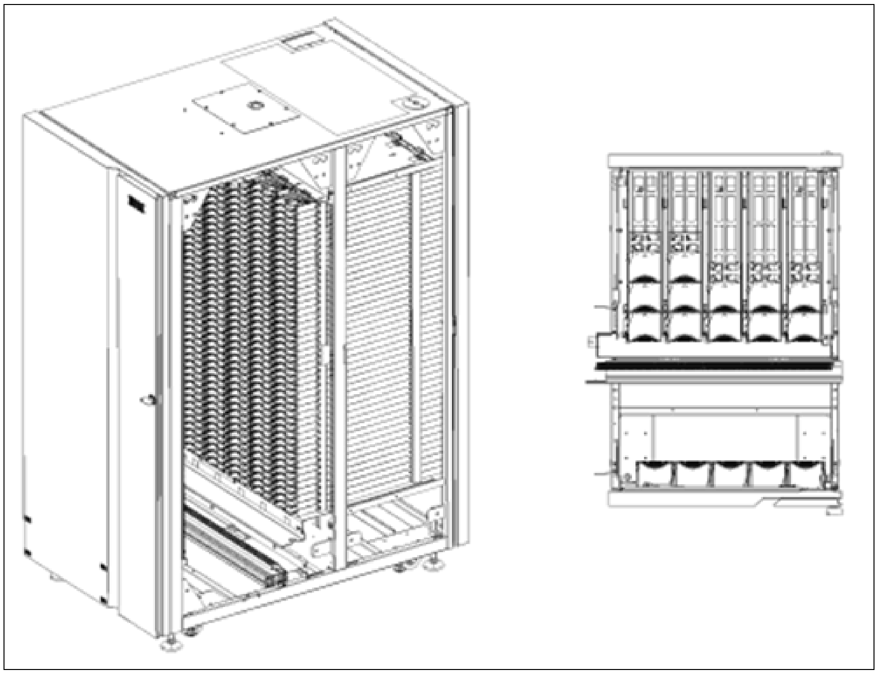

Figure 1-9 The HD frame (left) and top-down view of a row in an HD frame (right)

Models S24 and S54

The IBM TS3500 storage-only frame, HD1 models S24 and S54, can be attached to the TS4500 with the correct FC 1742 ordered.

These generation 1 HD frames can be redeployed into a TS4500 if they are installed to the right of the Lx5 frame. You must install FC 1742 before the frames can be added to a TS4500 library string. This feature code replaces the TS3500 cards to be supported on the TS4500, as shown in Figure 1-10.

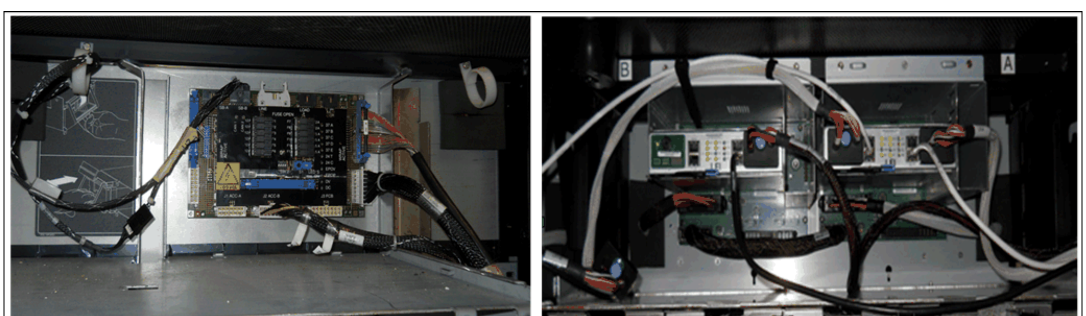

Figure 1-10 Changes to S24/S54 for TS4500 attachment

The TS3500 tape library models S24 and S54 frames are high-density (HD) version 1 storage-only expansion frames, which were attached to TS3500 tape libraries and frames.

The Model S24 expansion frame is for 3592 data cartridges. Up to 17 Model S24 expansion frames can be added to the right of the Lx5 frame of the TS4500 Model L25 base frame to increase 3592 cartridge storage. Each Model S24 frame supports up to 1,000 IBM 3592 cartridge slots.

The Model S54 expansion frame is for LTO data cartridges. Up to 17 Model S54 expansion frames can be added to the right of the Lx5 frame of the TS4500 tape library Model L55 base frame to increase LTO cartridge storage. Each Model S54 frame supports up to 1,320 LTO cartridge slots.

The HD1 models S24 and S54 can be added to any TS4500 expansion frame, if the expansion frame is added to the right of the Lx5 frame, up to a total of 18 expansion frames, including the Lx5 frame.

Note: The HD1 models S24 and S54 cannot be installed to the left of the Lx5 frame and cannot be installed as the rightmost frame in a dual accessor tape library.

HD1 and HD2 frames

All HD slots are black. However, the location of the cartridge retention latch differentiates LTO HD slots from 3592 HD slots. The cartridge retention latch is on the left side of LTO HD slots and on the right side of 3592 HD slots, as shown in Figure 1-11.

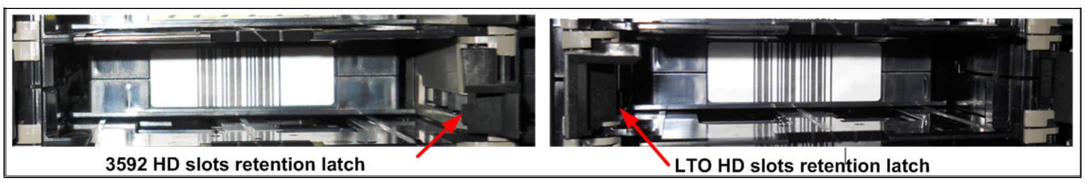

Figure 1-11 HD slot

Attention: The HD slots use a constant force spring to maintain forward pressure on the tape cartridges. Use caution when you insert or remove cartridges from the HD slots.

In an HD library, a standard inventory is a scan of tier 0 and tier 1. However, at times, it is necessary to inventory all tier. This operation takes more time because it requires moving the cartridges within an HD slot to scan each bar code. For all inventory operations, tier 2, and higher tier in an HD slot, are scanned only when one of the following changes occurs:

- A tier 1 cartridge bar code label was changed.
- Enough tier 1 bar code labels were changed in a column to warrant an inventory of the entire column of HD slots.
- Inventory of all tier is selected when you start a manual inventory from the TS4500 management GUI.

In HD frames, the cartridge accessor performs a shuffle operation to access the cartridges that are stored in tier 2 and higher. A *shuffle* is the process of moving cartridges in lower tier into the gripper, or other available slots, to access cartridges in higher tier (tier 2 or higher). To reduce shuffle operations and take advantage of repeated accesses of certain cartridges, the role of cartridge cache is assigned to all single-deep (tier 0) slots in an HD library.

To maintain efficient shuffle operations, the library uses load balancing to store cartridges across all HD slots in the library string. Therefore, all HD slots are filled to a minimum tier level until that tier is full across the library.

For the initial bulk load on a newly installed frame, insert cartridges into the deep slots, but leave the top two rows empty. The slots in the top two rows must be empty for the initial audit of the frame to start, and to enable the initial shuffle operation to proceed. The initial audit fills these slots, and then these slots are used like any other HD slot in subsequent library operations.

First-generation HD (HD1) frames can be installed to the right side of an Lx5 frame only.

Second-generation HD (HD2) frames provide the following enhancements:

- They can be installed in the leftmost library position (frame position 1).
- They offer drive-capable models that support up to 16 HD2-compatible tape drives when in frame position 2 or higher.

The position and configuration of the frame and the installation of CoD features affect the total available storage capacity of the S25 and S55 frames.

The available storage capacity based on possible frame model, position and configurations is listed in Table 1-8.

Table 1-8 Quantity of storage slots in the storage-only frames

Frame model	Licensed feature	F1 slots available	F2+ slots available
S25	Base	600	600
S25	HD CoD	800	1000
S24	Base	NA	600
S24	HD CoD	NA	1000
S55	Base	660	660
S55	HD CoD	1056	1320
S54	Base	NA	660
S54	HD CoD	NA	1320

1.2.4 TS4500 High Availability option

The TS4500 High Availability (HA) option provides a second accessor for redundancy and performance. This feature allows dual accessors in a dual active mode, and it features an integrated service bay to reduce service space.

The integrated service bays replace the HA frame and Service Bay B frames that are used on the TS3500. The integrated service bays allows a section of a frame to be used for servicing an accessor while the remaining portion of the frame is still available for active storage and drives. Accessor service must be performed through the side by opening the side doors of integrated service bays. Any HD2 frame can be an integrated service bay.

The second accessor, accessor B, is provided when you order a new Dx5 or Sx5 frame with FC 1442. The new Dx5 or Sx5 frame can be installed in any position, and it is included with the new B accessor. The accessor can be removed from the new frame and installed on the right side of the TS4500 if the new frame is installed on the left side.

Figure 1-12 shows the HA feature that is included in a new D25 frame.

Figure 1-12 D25 frame that was ordered with FC 1442

Restriction: The left or A side-integrated service bay (ISB) allows cartridges to be populated in storage columns 9 and 10, while drives can be populated in drive column 4 only.

The right or B side ISB allows cartridges to be populated in storage columns 1, 2, 3, and 4, while drives can be populated in drive columns 1 and 2 only. The I/O stations are accessible in the right integrated service bay.

An integrated service bay allows a minimum dual accessor system of only two frames.

The available storage for HA is shown in Figure 1-13.

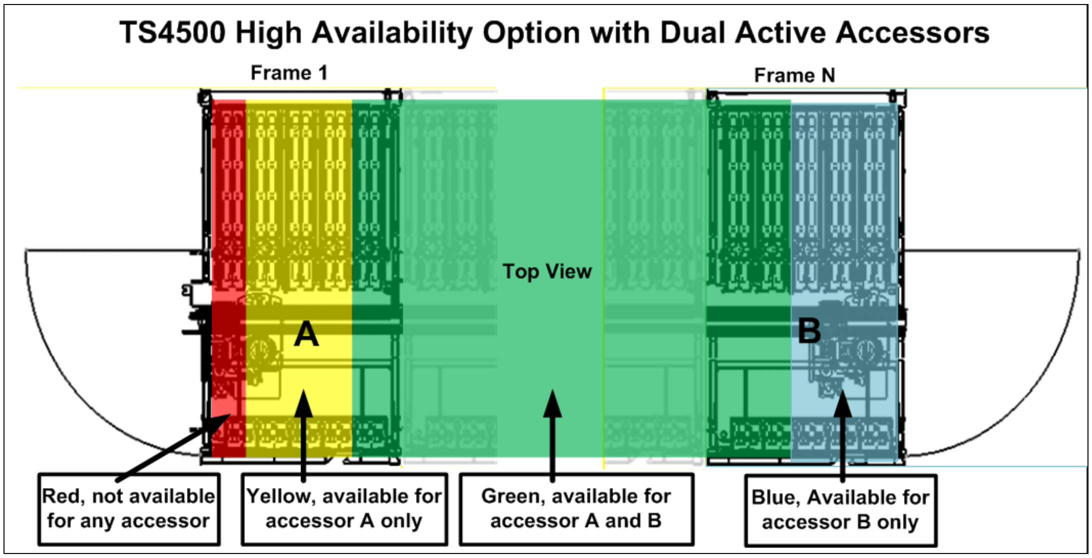

Figure 1-13 Available storage with an HA option

Note: I/O stations are not accessible in the left integrated service bay; therefore, an L25 or L55 frame cannot be used as a left integrated service bay. The frame that is included with the HA option can be installed on the left side (if required) because no dedicated service bays are on that side. The new accessor must be installed in the frame on the right side.

Elastic Capacity option

The TS4500 provides the Elastic Capacity option to completely eliminate inactive service space. The Elastic Capacity option can provide temporary relief for overflow conditions. These slots are referred to as tier T10 - T15 slots on LTO and T10-T14 on 3592. Standard storage slots are Tier 0 - 5 on LTO or Tier 0-4 on 3592.

With dual active accessors, certain storage slots are only available to a single accessor, and they are unavailable during accessor service. The ability to use these slots is optional, and they can be enabled or disabled on the management interface. The following modes are available for the Elastic Capacity option:

- Do not use
- Use for temporary overflow
- Use for maximum capacity

Do not use

In this mode, the TS4500 does not use the Elastic Capacity slots for media storage; therefore, all media is usable when only one accessor is available. Tier T10 - T15 slots are not used.

Use for temporary overflow

In this mode, cartridges are moved to elastic storage only if the library is 100% full, and if new inserts have no other destination choice. In this case, T10 and higher is used as temporary storage to handle the library overfill. Cartridges are returned to HA space after space becomes available manually by a user, or when the application mounts and unmounts the cartridges in the elastic capacity slots (T10 and higher).

Use for maximum capacity

Use this setting to store media cartridges in the limited access (Elastic Capacity) columns as normal storage. The use of these slots can be managed by using the following methods:

- Manually: A user can select a cartridge to destage to elastic storage by using the command-line interface (CLI). If the cartridge is already in elastic storage, no action is taken.
- Small Computer System Interface (SCSI): An application can use the HD Control field of the SCSI Move Medium command to specify that the move is an elastic storage destage.
- Periodic: When the dual-access area (the cartridge slots that both A and B accessors can reach) exceeds the usage threshold, the least recently used cartridges are moved into the Elastic Capacity area. The default usage threshold is 98%.

> **Note:** "Use for maximum capacity" is the default setting if the setting is not changed.

TS4500 capacity tool

The IBM Tape Library Slot Calculator helps calculate capacity and slot numbers for all IBM tape libraries, including TS3100, TS3200, TS3310, TS3500, and the TS4500. This calculation includes capacity with the different elastic capacity options.

This tool is available at this IBM Support web page.

Slot calculation is much more complex than totaling the numbers of slots in each frame. The slot calculator provides the available slot capacity that is based on any configuration. Of particular note is the "Elastic (Non-HA) capacity" utilization, which shows elastic usage with the current configuration. The does not use the elastic capacity until the number of CoD features reach a specific threshold.

The following examples show the effect on slot capacity when adding the HA feature (dual accessors). The examples show that overall capacity is available globally to the library, even though the CoD license keys are purchased against individual frame serial numbers.

Figure 1-14 shows an example of a 6-frame single accessor configuration with 3950 licensed slots available.

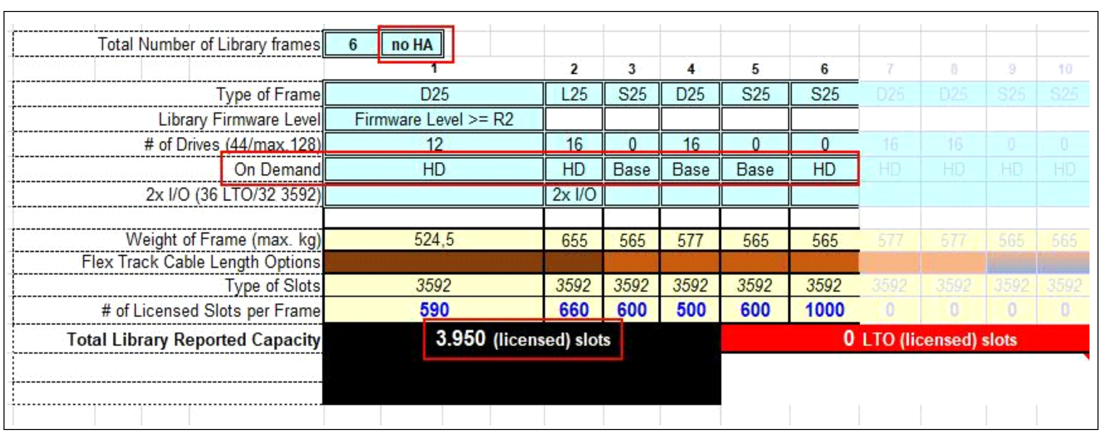

Total Number of Library frames	6	no HA								
	1	2	3	4	5	6	7	8	9	10
Type of Frame	D25	L25	S25	D25	S25	S25	D25	D25	S25	S25
Library Firmware Level	Firmware Level >= R2									
# of Drives (44/max 128)	12	16	0	16	0	0	16	16	0	0
On Demand	HD	HD	Base	Base	Base	HD	HD	HD	HD	HD
2x I/O (36 LTO/32 3592)		2x I/O								
Weight of Frame (max. kg)	524,5	655	565	577	565	565	577	577	565	565
Flex Track Cable Length Options										
Type of Slots	3592	3592	3592	3592	3592	3592	3592	3592	3592	3592
# of Licensed Slots per Frame	590	660	600	500	600	1000	0	0	0	0
Total Library Reported Capacity	3.950 (licensed) slots					0 LTO (licensed) slots				

Figure 1-14 6-frame configuration example: Single Accessor

Based on the same configuration, if we add the HA feature to make it a dual accessor library, we can see the effect that this addition has on slot numbers.

> **Note:** Frame 6 now includes the service bay for the second accessor, which means that access to some columns and therefore storage slots are restricted. However, although the library is not fully licensed with CoD features, the firmware compensates for the loss of the slots that is caused by the second accessor installation.

In Figure 1-15, we can see that adding the dual accessor (HA) has no effect on the number of licensed storage slots, and that the elastic capacity is not used.

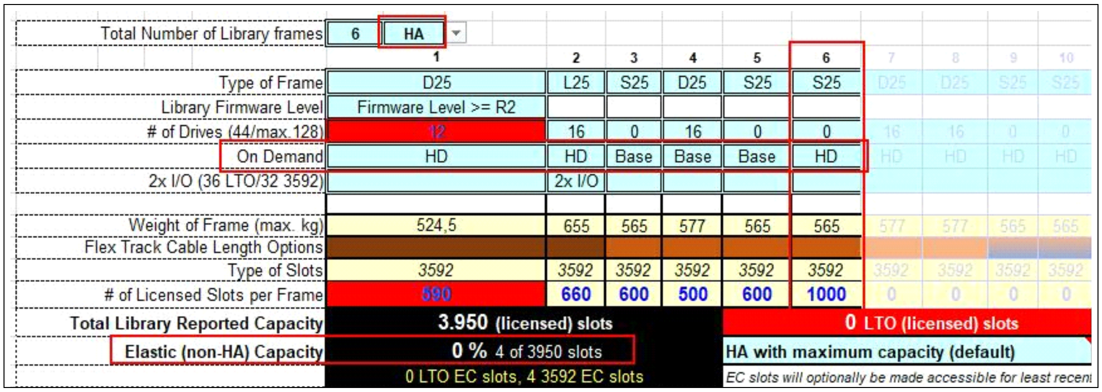

Total Number of Library frames	6 HA									
	1	2	3	4	5	6	7	8	9	10
Type of Frame	D25	L25	S25	D25	S25	S25	D25	D25	S25	S25
Library Firmware Level	Firmware Level >= R2									
# of Drives (44/max.128)	12	16	0	16	0	0	16	16	0	0
On Demand	HD	HD	Base	Base	Base	HD	HD	HD	HD	HD
2x I/O (36 LTO/32 3592)		2x I/O								
Weight of Frame (max. kg)	524,5	655	565	577	565	565	577	577	565	565
Flex Track Cable Length Options										
Type of Slots	3592	3592	3592	3592	3592	3592	3592	3592	3592	3592
# of Licensed Slots per Frame	590	660	600	500	600	1000	0	0	0	0
Total Library Reported Capacity	3.950 (licensed) slots				0 LTO (licensed) slots					
Elastic (non-HA) Capacity	0 % 4 of 3950 slots				HA with maximum capacity (default)					
	0 LTO EC slots, 4 3592 EC slots				EC slots will optionally be made accessible for least recent					

Figure 1-15 6-frame configuration example: Dual Accessor

1.2.5 Mainframe-ready

From FW release 3, the TS4500 is supported on IBM z/OS systems with an attached TS7700. The TS7700 requires FW release 4 and higher to attach to a TS4500.

To provide this support, the TS4500 must have two integrated 16 Gb Fibre Channel switches, which can be installed in the bottom of a L25 or D25 frame, as shown in Figure 1-16.

Figure 1-16 Integrated TS7700 back-end switches

If a migration of the 16 Gb Fibre Channel Switch is from a TS3500 TopRack (TR1) to be installed on a TS4500, the following options are available:

- TopRack can be reinstalled with the switches on the TS4500; the existing TS3500 TR1 can be moved to the TS4500 (see Figure 1-19 on page 30).
- Order only FC 4879 (TS7700 BE Switch Mounting Hardware), which provides the switch mounting kit without switches. This option is provided as a field MES if the TS4500 is already installed.

The TS4500 management GUI supports the preset TS7700 logical library and the use of an external TSSC/IMC.

> **Note:** The integrated Fibre Channel switches do not require extra power feeds to the TS4500. The integrated Fibre Channel switches use bifurcated power cords that are provided with the mounting kit hardware for the switches. These power cords connect internally to the existing TS4500 Power Distribution Units.

1.2.6 External TSSC/IMC

The TS4500 supports the use of an external TSSC/IMC so that a single TSSC provides the Call Home capability for several TS7700, TS4500, or TS3500 devices on the same site, as shown in Figure 1-17.

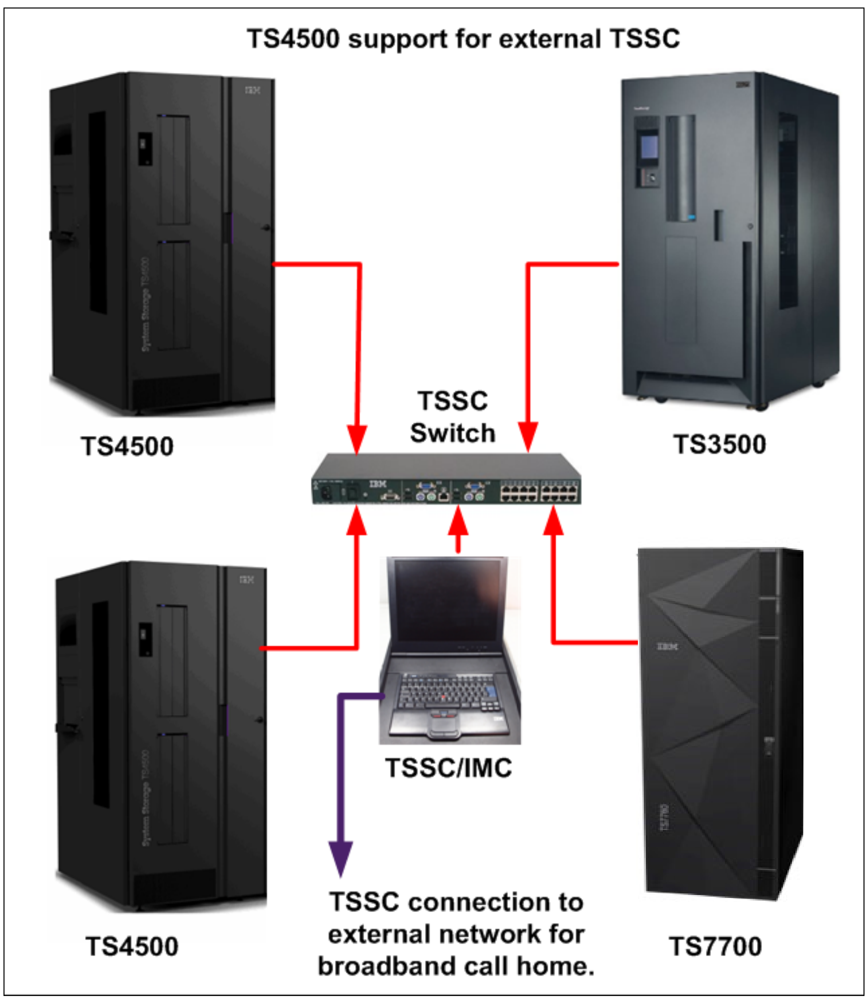

Figure 1-17 External TSSC server

> **Note:** The external TSSC can also be a TSSC/IMC inside another TS4500. This configuration requires that you install FC 2704 on the Lx5 frame. FC 2704 provides a 26-port switch to allow connection to up to 24 extra devices to share the TSSC/IMC.
>
> Only one TSSC/IMC can be configured for IBM Call Home: the external TSSC/IMC or by using an internal IMC.

1.2.7 TS4500 tape library top rack frame TR1 and TR2

The TS4500 top racks provide extra rack space on any frame in a library without requiring more floor space. They also simplify cabling by providing extra rack space above the library for power distribution units, Fibre Channel switches, tape data movers, or IBM Linear Tape File System (LTFS) nodes.

Both Top Rack models (TR1 and TR2) are installed in the field by an IBM service representative or service partners on one or more frames. The top racks, and any components that are housed in the racks, are supported and serviced independently of the TS4500 tape library.

FC 1750, top rack end covers, is required for the left and right ends of one or more adjacent top racks. This feature is required for only the first top rack that is ordered when multiple top racks are ordered for adjacent frames.

3584 Model TR1

The optional top rack, 3584 Model TR1 (see Figure 1-18), provides an extra 10U of rack space on any frame. The components that are placed into the top rack should not exceed 30 lbs per U, which is a maximum of 300 lbs for the TR1.

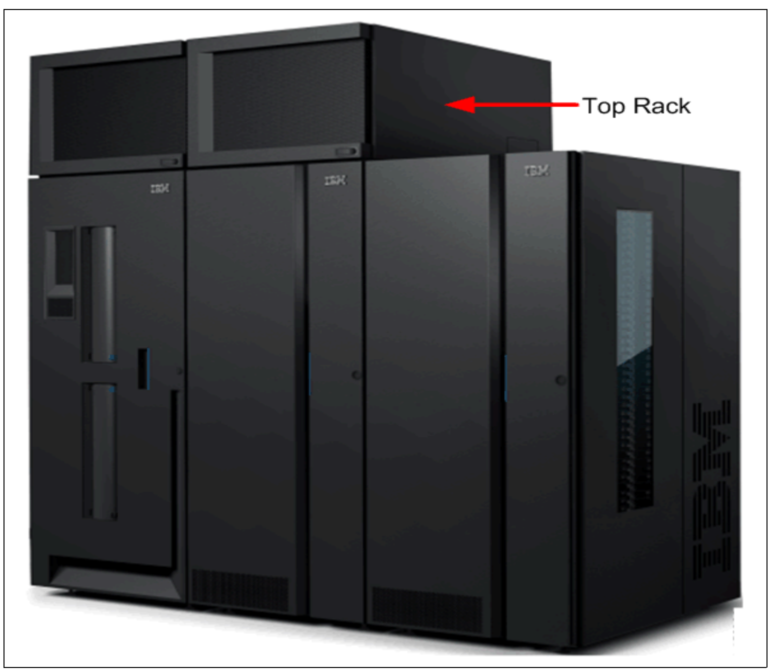

Figure 1-18 Tape library with an installed top rack TR1

FC 1751 power distribution unit (PDU), or 1752 Enhanced PDU, optionally can be ordered. Up to two PDUs can be ordered for the TR1. The first PDU does not use any of the 10U rack space. A second PDU (for redundancy) uses 1U of rack space.

One power cord feature, 9954 - 9959 or 9966, is required for each 1751 feature that is ordered. One power cord feature, 9954 - 9958 or 9948, is required for each 1752 feature that is ordered.

3584 Model TR2

The TR2 provides an extra 5U of rack space on any frame. The lower overall size of Model TR2 compared to Model TR1 enables you to install TR2 where overhead space prevents TR1 from being installed. The components that are placed into the top rack should not exceed 30 lbs per U, which is a maximum of 150 lbs for the TR2. FCs 1755 (Front Door) and 1756 (Rear Door) are optional.

Rack configurations

The following rack configurations are available:

- Standard 19-inch Rack (FC 1754)

  The TR2 is configured in conformance with industry standard 19-inch racks. This configuration provides 5U of 19-inch rack space.

- Olympus Rack (Feature code 1753)

  The TR2 rack is configured in conformance with the Open Compute Project Olympus Rack Specification.

FC 1752, Enhanced PDU, optionally can be ordered for the TR2. Each PDU uses 1U of rack space. Up to two of FC 1752 can be ordered. Each Enhanced PDU provides six C13 outlets and three C19 outlets.

Power cords

Consider the following points:

- For single phase input power, each PDU order you must order one power cord FCs 9954 - 9958.
- For 3-phase (wye) input power, each PDU that you order must order one power cord FC 9948.

Figure 1-19 shows the top rack without covers and devices installed.

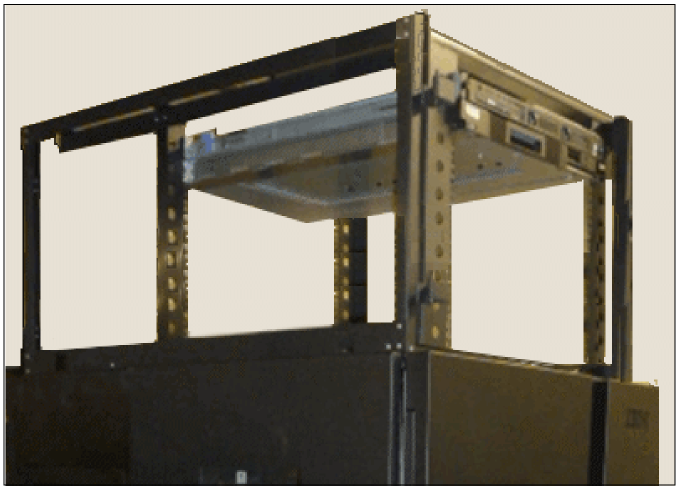

Figure 1-19 Top rack without covers and devices

Important: The top rack is treated as an independent rack space, and it is not tied to the service or support of the tape library.

1.3 TS4500 tape library components

The TS4500 tape library consists of one or more frames, which include more components that supply power to the library, installed tape drives, and components for handling and storing tape cartridges.

Each available component in the front and side of the frame is shown in Figure 1-20.

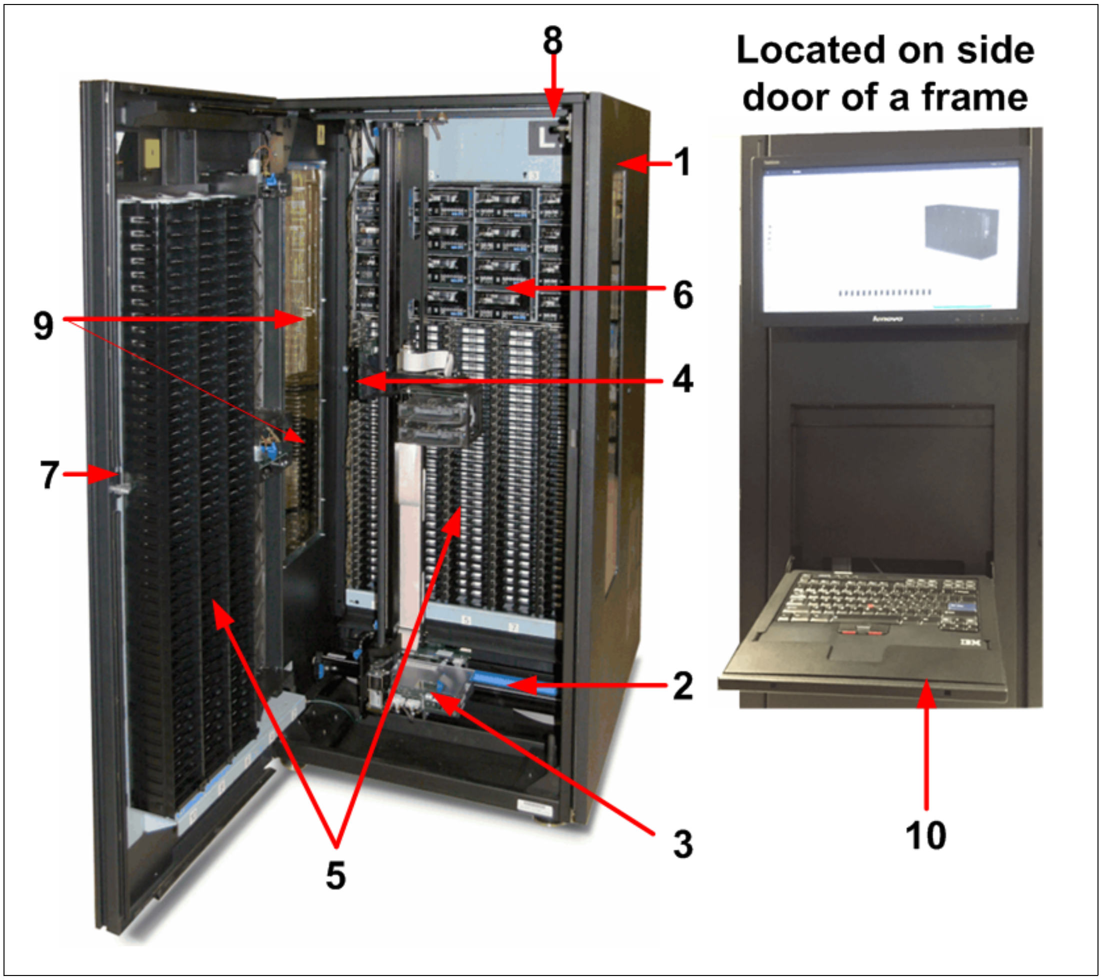

Figure 1-20 TS4500 components that are available from the front of the frame

The components that are shown in Figure 1-20 on page 31 are listed in Table 1-9.

Table 1-9 Inside the front of the TS4500

Number	Component
1	Library frames
2	Rail system
3	Cartridge accessor
4	Accessor controller
5	Cartridge storage slots
6	IBM LTO or 3592 tape drives
7	Front door
8	Door safety switch
9	I/O stations
10	Integrated management console

Accessor B is in the right-side frame. Accessor B is functionally the same as accessor A, as shown in Figure 1-21.

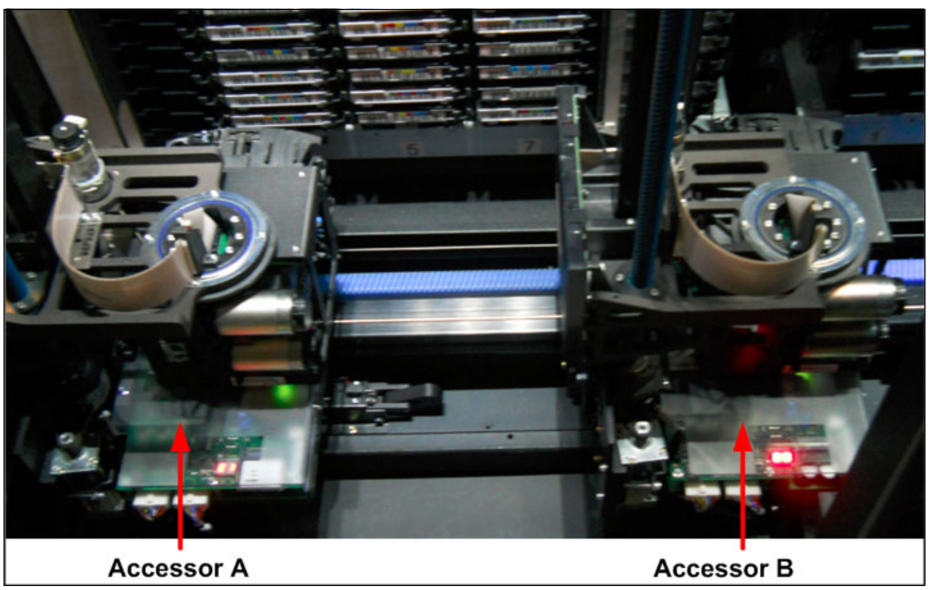

Figure 1-21 Dual accessor

The location of each component that is available in the top and rear of the frame is shown in Figure 1-22.

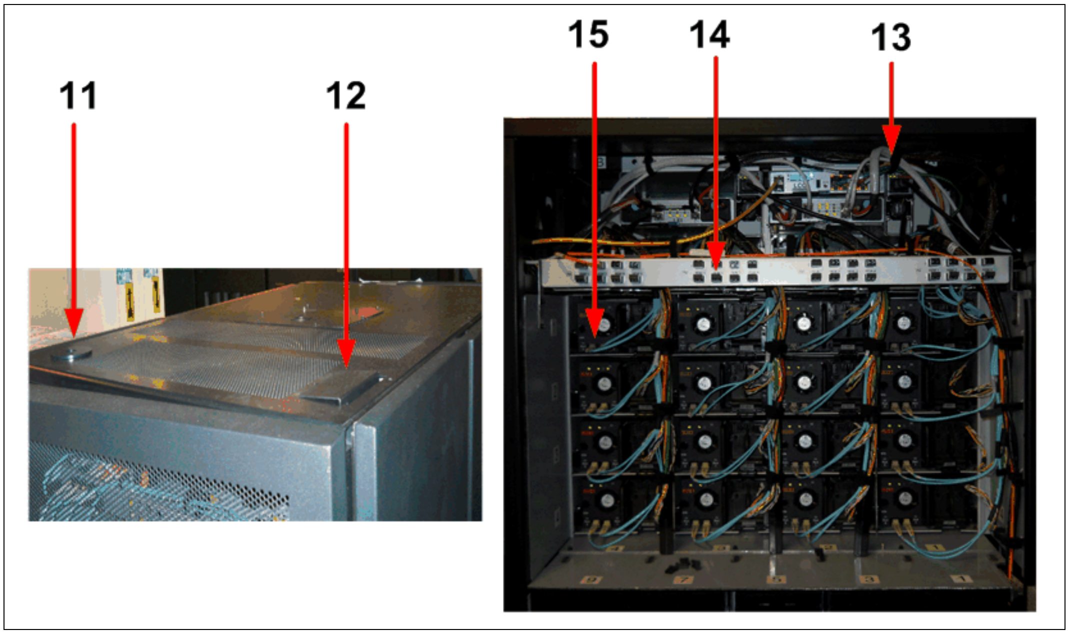

Figure 1-22 Top and rear of the TS4500

The components that are shown in Figure 1-22 are listed in Table 1-10.

Table 1-10 Top and rear components of the TS4500

Number	Component
11	Power cable hole
12	Fibre Channel cable hole
13	Frame control assembly
14	Patch panel
15	Rear of IBM LTO or 3592 tape drives

The location of each component that is available on the front door is shown in Figure 1-23.

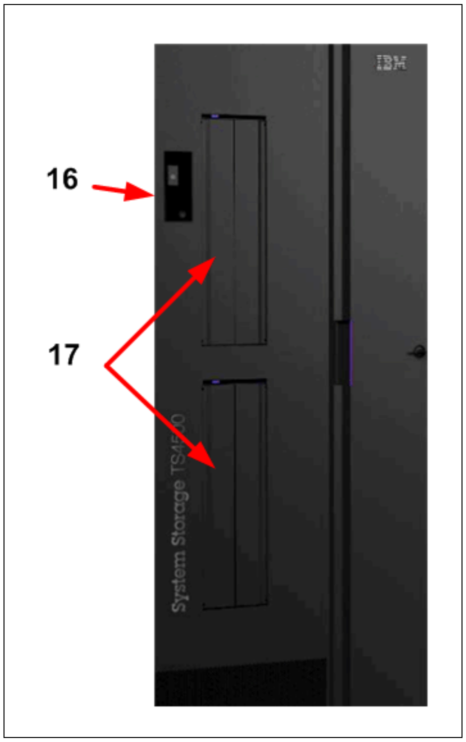

Figure 1-23 Front door

The components that are shown in Figure 1-23 are listed in Table 1-11.

Table 1-11 Front door components of the TS4500

Number	Component
16	Display panel
17	I/O stations

Library frames

The base frame (Lx5 models) and the expansion frames (Dx5, Sx5, and Sx5 models) are the building blocks of the library. Each frame contains a rail system, high-density cartridge storage slots, and internal LED lighting. The Lx5 and Dx5 frames also contain slots for up to 16 tape drives.

Rail system and track cable

The cartridge accessor moves through the TS4500 tape library on a rail assembly. The system consists primarily of a main rail assembly and a support rail, and a trough for the power and control cable. The main rail assembly includes a main bearing way with a rack gear. Its support rail is an L-shaped rail that runs along the top of the frames and provides smooth transport for the cartridge accessor.

TS4500 has a newly designed flex track cable and guide. This new flex track cable and guide was designed to reduce cost and outage time when you add a frame (see Figure 1-24).

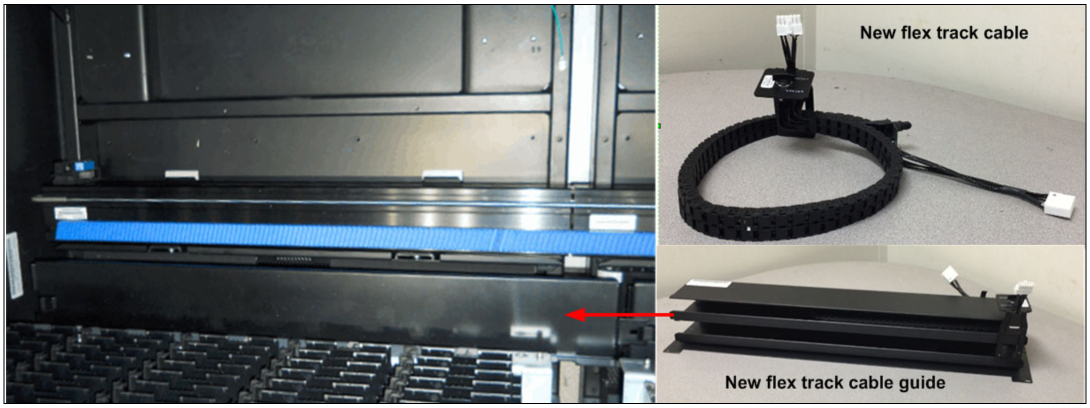

Figure 1-24 Rail system and track cable

The new style of cable requires no tools to install or replace. It is stacked in a new guide that is designed with two-chambers, which separate signal and power wires. Accessor A uses the top guide and accessor B uses the bottom guide, as shown in Figure 1-25.

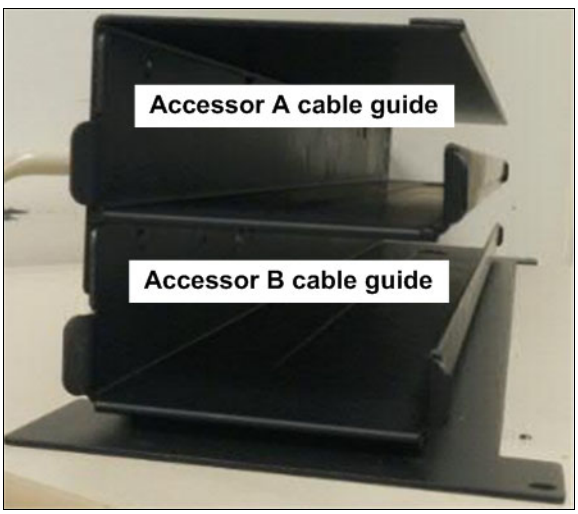

Figure 1-25 Track cable guide

These cables plug into the same frame. Depending on the library size, chose the suitable cable length and install in the correct frame, as listed in Table 1-12.

Table 1-12 Flex track installation frames

Flex track cable size	**Installations in frame**
1 - 2 frames	1
3 - 6 frames	3
7 - 14 frames	7
9 - 18 frames	9

Important: When you add the new HA feature to a TS4500, you must order FC 2071 or FC 2072 for each frame in addition to ordering the correct length flex track cable.

Cartridge accessor

The *cartridge accessor* moves cartridges between the storage slots, tape drives, and the I/O station of the TS4500 tape library. If the HA option is installed, two accessors exist: accessor A and accessor B. Functionally, accessor A and accessor B are identical. They have the same hardware components, except for a bottom bumper and the location of the X home sensor moving to the right side of accessor B.

The accessors consist of the components that are described next.

X-axis and Y-axis motion assemblies

This group of parts includes a controller (circuit board) for the Controller Area Network interface, servo motor, pinion drive gear, and lead screw. These assemblies provide the motive force to move the accessor side to side (on the X-axis) and up and down (on the Y-axis). The controller part of this assembly is referred to as the *XY controller*, as shown in Figure 1-26.

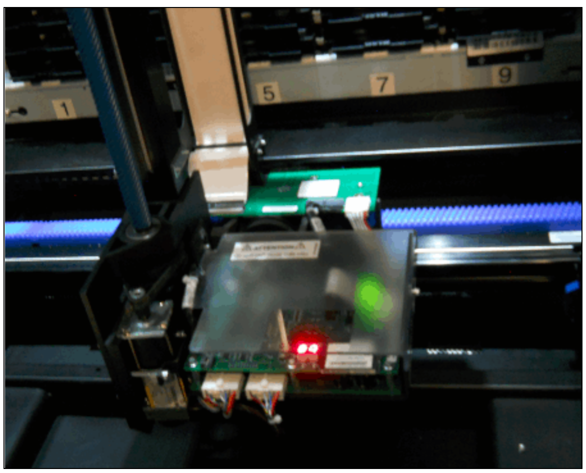

Figure 1-26 Motion assembly

Pivot assembly

This group of parts provides a mounting platform for the gripper mechanism and the bar code reader. This assembly can rotate 180° around the vertical axis, as shown in Figure 1-27.

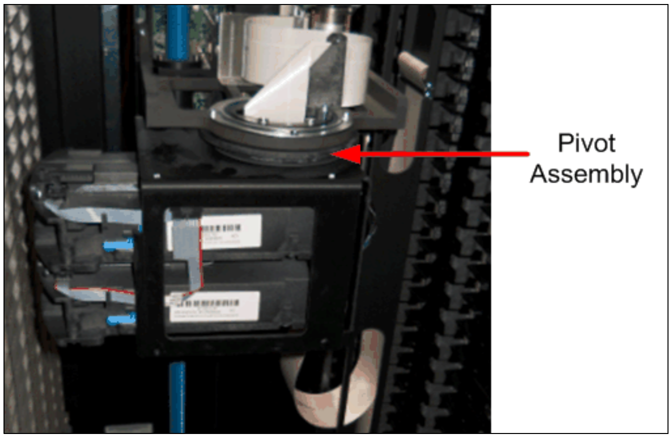

Figure 1-27 Pivot assembly

Optimized dual gripper

This electromechanical device (which is mounted on the pivot assembly) gets or puts cartridges from or to a storage slot, tape drive, or I/O station. The gripper is independently controlled, and it can grip a single cartridge. Two grippers are on the pivot assembly (Gripper 1 and Gripper 2). The grippers are in the dual-gripper transport mechanism, as shown in Figure 1-28.

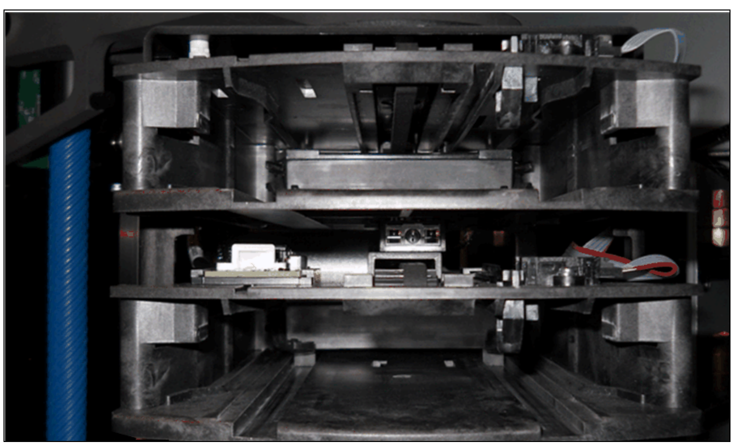

Figure 1-28 Dual gripper

Bar code reader and calibration assembly

This assembly has a dual purpose: It reads the bar code on a label that is affixed to a cartridge or to the rear of empty storage slots, and calibrates the frame after the installation or hardware change. The bar code reader/calibration sensor is mounted on the bottom of the top gripper, and it can be seen between the two grippers when viewed from the front. It is used for frame calibration during inventories, audits, insertions, and inventory updates (a process that is invoked each time that you open a door).

An Auxiliary Lighting Element (ALE) assists the bar code reader. The ALE provides light to assist the bar code reader to scan labels correctly. The ALE connects to the top of the bottom gripper.

The inventory update determines whether cartridges are added to or removed from the library, or moved within the library. The bar code assembly, calibration assembly reader, and ALE are shown in Figure 1-29.

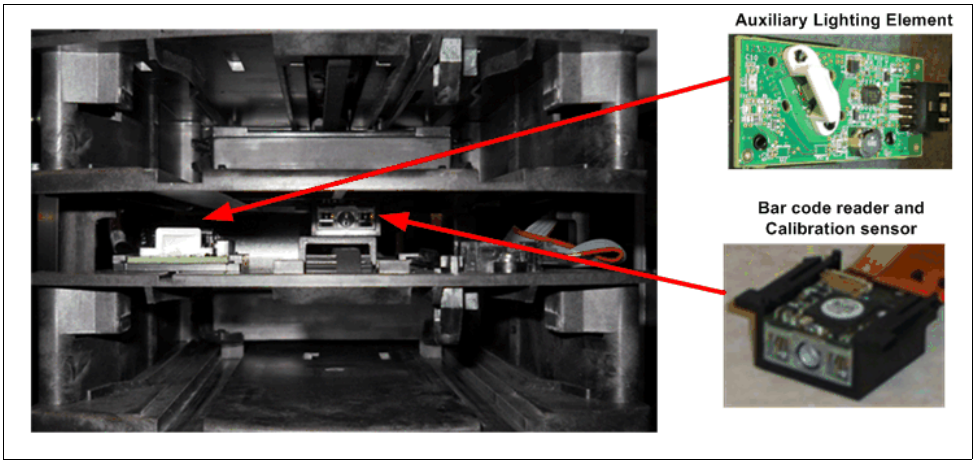

Figure 1-29 Bar code reader, calibration sensor, and ALE

Accessor controller

This controller is a circuit board that facilitates all accessor motion requests, such as calibrations, moves, and inventory updates. This controller is on the side of the accessor assembly, as shown in Figure 1-30.

Figure 1-30 Accessor controller

Cartridge storage slots

All frames contain single-deep cells that are mounted on the door of the frame. Each cell stores one tape cartridge. High-density cells are mounted on the inside wall of each frame. These high-density cells each store four (3592) or five (LTO) tape cartridges. Individual frames do not support mixed media (a combination of 3592 and LTO tape cartridges). However, mixed media is supported within the TS4500 tape library.

Release 2 introduces new single-deep cells for both LTO and 3592 cartridges that are on the door side of HD2 frames. A chevron fiducial, which is similar to the chevron fiducial that is on the deep cells for individual cell calibration, is included in the new single-deep cells. Ribs are included to reduce cell wear, and they are a new feature of the single-deep cells, as shown in Figure 1-31.

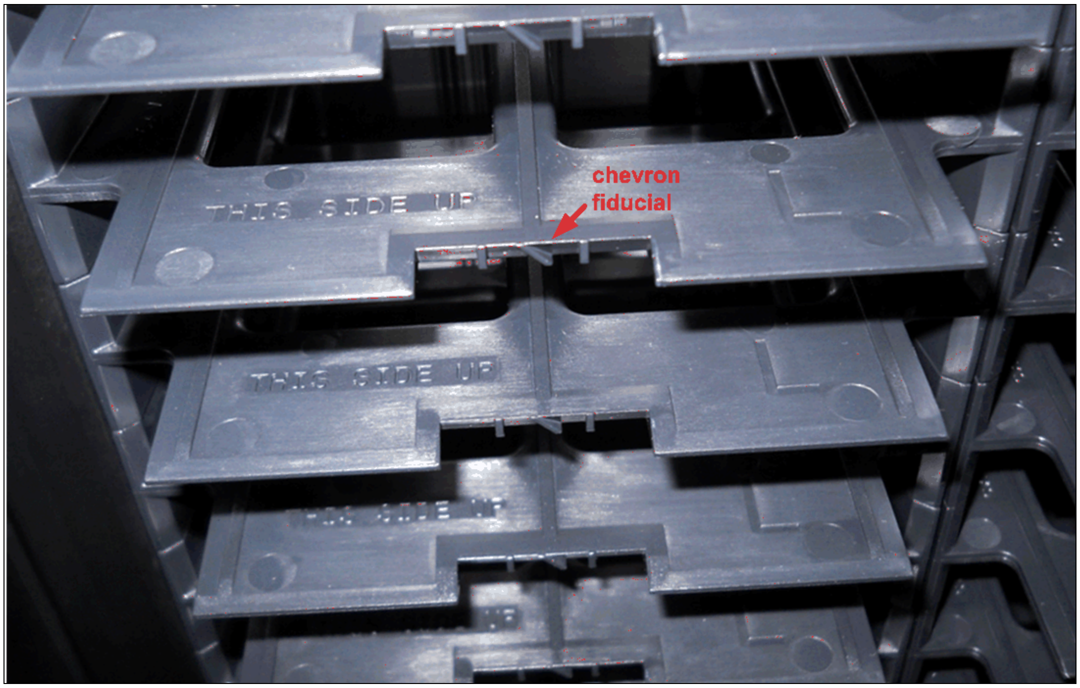

Figure 1-31 New door side single-deep cell slots

IBM LTO or 3592 tape drives

Drive frames can contain one or more units that are mounted in the frame. The TS4500 tape library supports LTO and 3592 tape drives. The HD2 frames of the TS4500 tape library support HD2-compatible models of the TS1160, TS1155, TS1150, TS1140, LTO-9, 8, 7, 6, and 5 tape drives.

Up to 12 drives can be installed in an Lx5 or Dx5 frame that is in frame position 1 (the leftmost frame) of the library. Up to 16 drives can be installed in each Lx5 or Dx5 frame that is in frame position 2 or higher. Within a HA configuration, the maximum number of drives that can be installed in the first frame is four and in the last frame is eight, if the frame is a Lx5 or Dx5 frame.

LTO and 3592 tape drives cannot be mixed in the same frame, but the LTO and 3592 frames can be mixed in the same library. You can identify a drive by inspecting the label at the rear of the drive canister. For more information about these drives, see Chapter 2, “TS4500 Ultrium Linear Tape-Open and 3592 tape drives” on page 77.

Note: No drives are installed in the storage-only frames (models S25, S24, S55, and S54).

The supported tape drives are listed in Table 1-13.

Table 1-13 Tape drives that are supported by the TS4500 tape library

Type of drive	Speed of connection	Native data rate	Native capacity read/write	Other information
IBM LTO Ultrium 9	8 Gbps Fibre for model F9C and F9S 12 Gbps SAS for model S9C	400 MBps	18 TB (16.37 TiB)	Known as the LTO-9 tape drive, Model 3588 F9C, F9S and S9C, TS1090
IBM LTO Ultrium 8	8 Gbps Fibre	360 MBps	12 TB (10.91 TiB)	Known as the LTO-8 tape drive, Model 3588 F8C, TS1080
IBM LTO Ultrium 7	8 Gbps Fibre	300 MBps	6 TB (5.46 TiB)	Known as the LTO-7 tape drive, Model 3588 F7C, TS1070
IBM LTO Ultrium 6	8 Gbps Fibre	160 MBps	2.5 TB (2.27 TiB)	Known as the LTO-6 tape drive, Model 3588 F6C, TS1060
IBM LTO Ultrium 5	8 Gbps Fibre	140 MBps	1.5 TB (1.36 TiB)	Known as the LTO-5 tape drive, Model 3588 F5C, TS1050
IBM TS1160	16 Gbps Fibre for model 60F 10 or 25 Gbps Optical Ethernet for model 60E 12 Gbps SAS for model 60S	400 MBps	► 900 GB (.82 TiB) with JK ► 5 TB (3.63 TiB) with JM ► 7 TB (6.37 TiB) with JC/JY ► 15 TB (13.64 TiB) with JD/JZ ► 20 TB (18.19 TiB) with JE/JV	Known as the 3592 60E, 3592 60F, or 3592 60S tape drive
IBM TS1155	8 Gbps Fibre for model 55F 10 Gbps Optical Ethernet for model 55E	360 MBps	► 900 GB (.82 TiB) with JK ► 3 TB (2.73 TiB) with JL ► 7 TB (6.37 TiB) with JC/JY ► 15 TB (13.64 TiB) with JD/JZ	Known as the 3592 55E or 3592 55F tape drive

Type of drive	Speed of connection	Native data rate	Native capacity read/write	Other information
IBM TS1150	8 Gbps Fibre	360 MBps	► 900 GB (.82 TiB) with JK ► 2 TB (1.82 TiB) with JL ► 7 TB (6.37 TiB) with JC/JY ► 10 TB (9.1 TiB) with JD/JZ c	Known as the 3592 EH8 tape drive
IBM TS1140	8 Gbps Fibre	250 MBps	► 500 GB (.48 TiB) with JK ► 1.6 TB (1.46 TiB) with JB/JX ► 4 TB (3.6 TiB) with JC/JY	Known as the 3592 EH7 tape drive

Figure 1-32 shows the TS4500 drive bay.

Figure 1-32 TS4500 drive bay

Front door

The front door contains single cartridge storage slots on the inside of the door that are referred to as *tier 0 slots*. Two I/O stations are installed on the front door of the base frame. Optionally, two extra I/O stations can be installed on the front door of any Dx5 frame. The library's front door has a key lock. The key lock is the same for every front door, and the keys are included with the library. The front door is shown in Figure 1-23 on page 34.

Note: The side doors of the base frames (Lx5) also have a key lock, which is the same key lock that is used for the front door.

Door safety switch

This switch ensures that power to the cartridge accessor is switched off whenever the front door or side door is opened. This safety component ensures that accessor movement is stopped while the front door is open. Figure 1-33 shows the front door safety switch.

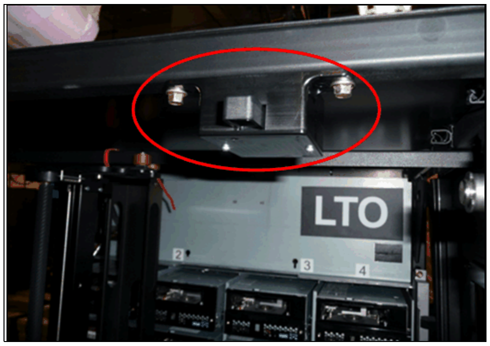

Figure 1-33 Front door safety switch

Side door and accessor service access

To service a single accessor concurrently, the accessor that requires service must be placed into service mode from the management GUI. After the accessor parks, and the side door is open, the service bay switch (SBS) activates. This sequence puts the SBS into the up position. Figure 1-34 shows the SBS mechanism.

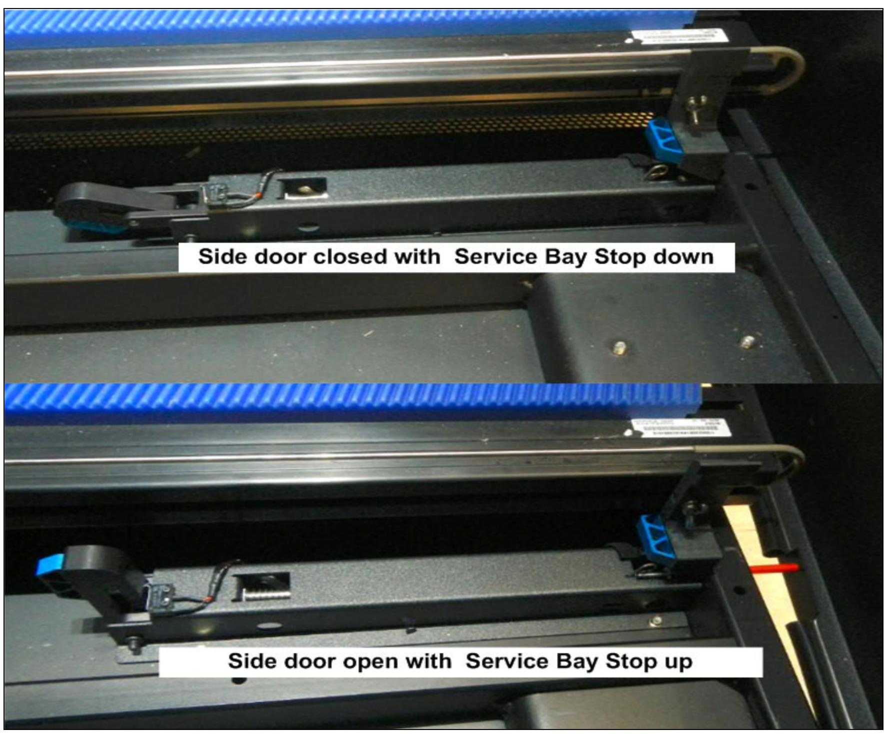

Figure 1-34 Service bay switch mechanism

Figure 1-35 shows an accessor in service with side door open and end stop up. The second accessor is working and it is prevented from moving into the service area by the end stop.

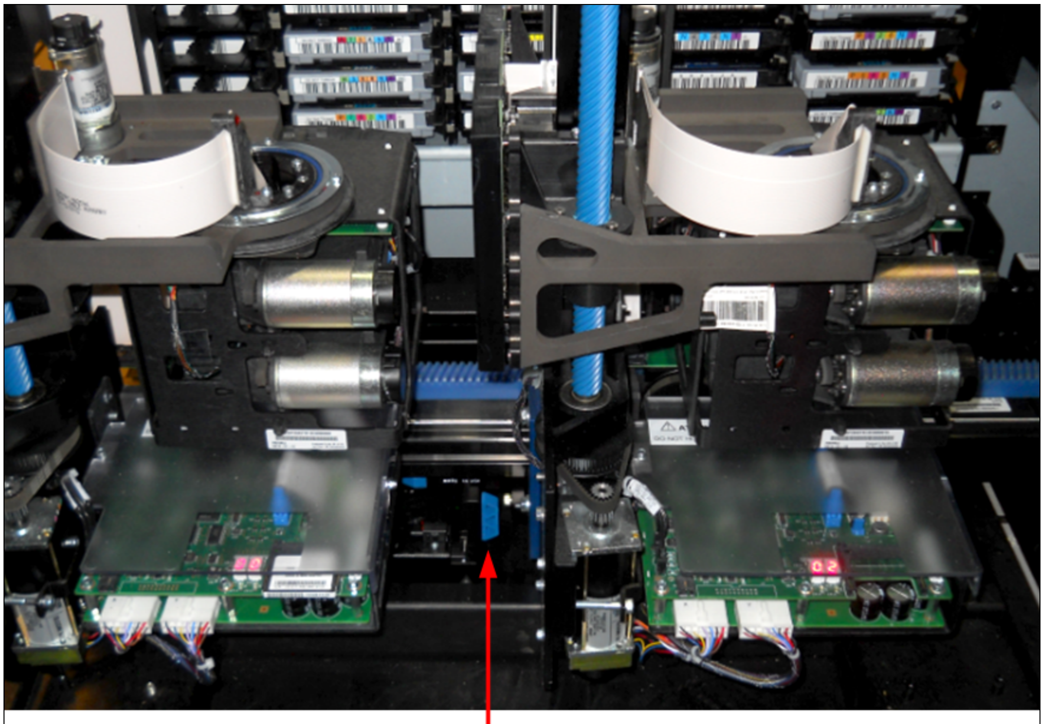

Accessor in service with side door open and Service Bay Stop in place

Figure 1-35 Accessor in service

To service the accessor, it must be removed from the TS4500 by using the side door. This action is a concurrent action with the second accessor that is used for move commands. This process is simple because the end stops and track cable can be removed without any tools.

I/O stations

Two I/O stations on the front door of the base frame enable the insertion or removal of tape cartridges without requiring the library to reinventory the frame. Optionally, two more I/O stations can be installed on any Dx5 expansion frame. The I/O station is universal, and either LTO or 3592 magazines can be installed in it.

Cartridges can be inserted or removed by using the I/O stations while the TS4500 tape library performs other operations without requiring an inventory.

The TS4500 tape library base frames (models L25 and L55) come with two I/O stations. Each I/O station houses a cartridge magazine so that individual cartridges can be handled independently of the tape library. Consider the following points:

- A cartridge magazine for LTO can hold up to 18 cartridges.
- A cartridge magazine for 3592 can hold up to 16 cartridges.

On libraries with LTO and 3592 frames, the first Dx5 frame (different from the Lx5) can have I/O stations that are included with magazines for the different drive type. The LTO and 3592 magazines can be installed on any I/O station.

Magazines for different media types can also be ordered by using FC 1628 for LTO and FC 1629 for 3592.

The handles on the cartridge magazine are used to insert and remove the magazine, or to carry it during transport. The magazine safety lock, as shown in Figure 1-36, retains cartridges in their slots and prevents them from falling out while the magazine is transported.

Figure 1-36 Tape cartridge magazine safety lock

When a magazine is removed from an I/O station, it is necessary to engage the safety lock until the magazine is placed on an accessible surface. It is then necessary to unlock the magazine to insert or remove cartridges. Attempting to insert or remove cartridges while the magazine safety lock is engaged might damage the magazine or the cartridges.

After a magazine is reinserted into an I/O station, it is necessary to unlock the magazine so that the accessor can retrieve cartridges. The I/O doors do not close correctly if the magazine is not unlocked.

Figure 1-37 shows a full 3592, a full LTO magazine, and an empty 3592 magazine.

Figure 1-37 I/O magazines

Two more I/O stations can be installed in any Dx5 expansion frame by ordering FC 1652. This feature installs two I/O stations in one expansion frame. Each additional pair of I/O stations increases the maximum insert/eject throughput for the library. The maximum cartridge capacity for expansion frames with two I/O stations is reduced by 80 cartridges for Model D25 and by 88 cartridges for Model D55. Storage-only frames (models Sx5) do not support I/O stations.

You can remotely use the I/O station action menu, which is available from the System page of the TS4500 management GUI, to open and close the I/O station doors. At the library, you can press the eject button (which is numbered 1 in Figure 1-38 on page 46) to open and close the doors. When the doors are open, it is possible to manually remove and replace the cartridge magazine to insert or remove cartridges.

When the doors are closed, the cartridge accessor can access the cartridges. The lock status LED that is next to the eject button (numbered 2 in Figure 1-38 on page 46) indicates that the I/O station is locked because the accessor can insert or remove cartridges. Do not attempt to open the I/O station when the lock status indicator is illuminated.

Figure 1-38 shows the top I/O station and panels, which control and show the state of the I/O station.

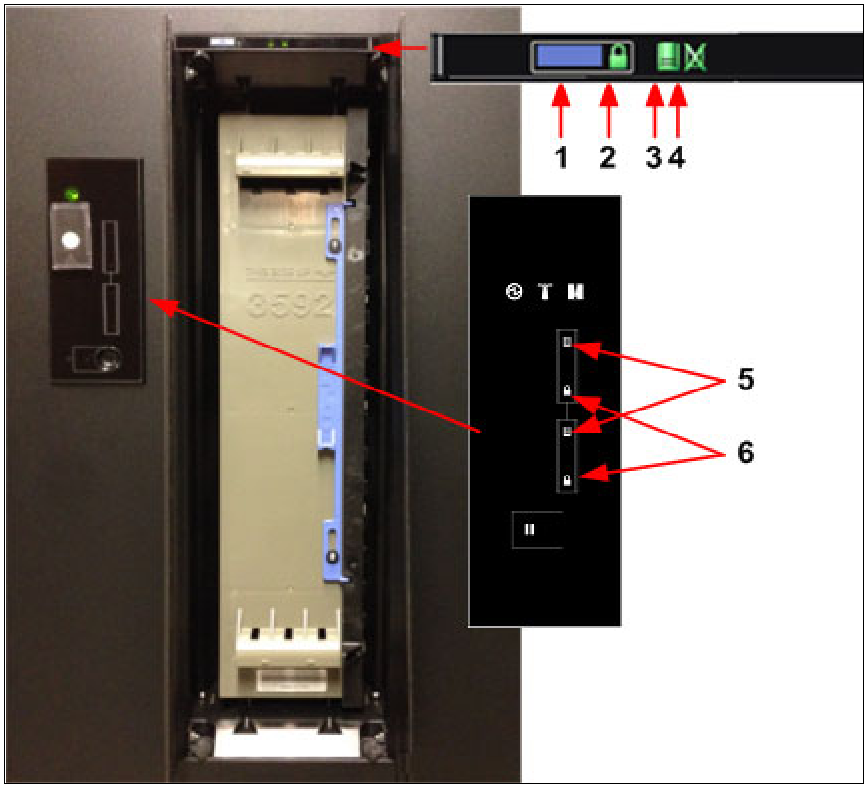

Figure 1-38 I/O station LEDs and panels

Important: Do not tilt the magazine during installation. Not tiling the magazine avoids pushing the magazine through the I/O station and obstructing the accessor.

The I/O station controls and LEDs are listed in Table 1-14.

Important: Use only the eject button (which is numbered 1 in Figure 1-38 on page 46) or the management GUI to open and close the I/O station doors. Do not attempt to open or close the doors manually.

Table 1-14 I/O station (numbers correspond to Figure 1-38 on page 46)

Number	Function
1	Eject button.
2	► Open/close state. ► Off: Normal state. ► Flashing: Transitional state. Whenever the button is pushed, it flashes and goes off when the doors open, and it flashes and comes back on as a solid light if it is locked by a “code”. ► On solid: Locked.
3	Unload is required, or cartridges are present. ► Off: No cartridges are present. ► Flashing: The I/O station is full so an intervention is required. ► On: Several cartridges are present.
4	No magazine is present. On solid: No magazine is present.
5	I/O fullness indicator (green)
6	I/O lock indicator (green)

If the I/O station is obstructed, the doors automatically reopen.

Each I/O station slot features a unique address to indicate its physical location. The I/O station slot address consists of two values: a frame number and a row number.

After you close the I/O station doors, the library automatically moves the cartridges into storage slots. How each cartridge is assigned to a logical library depends on the configured VOLSER ranges. For more information, see Figure 1-39 on page 48. The state of the cartridge is shown in bold.

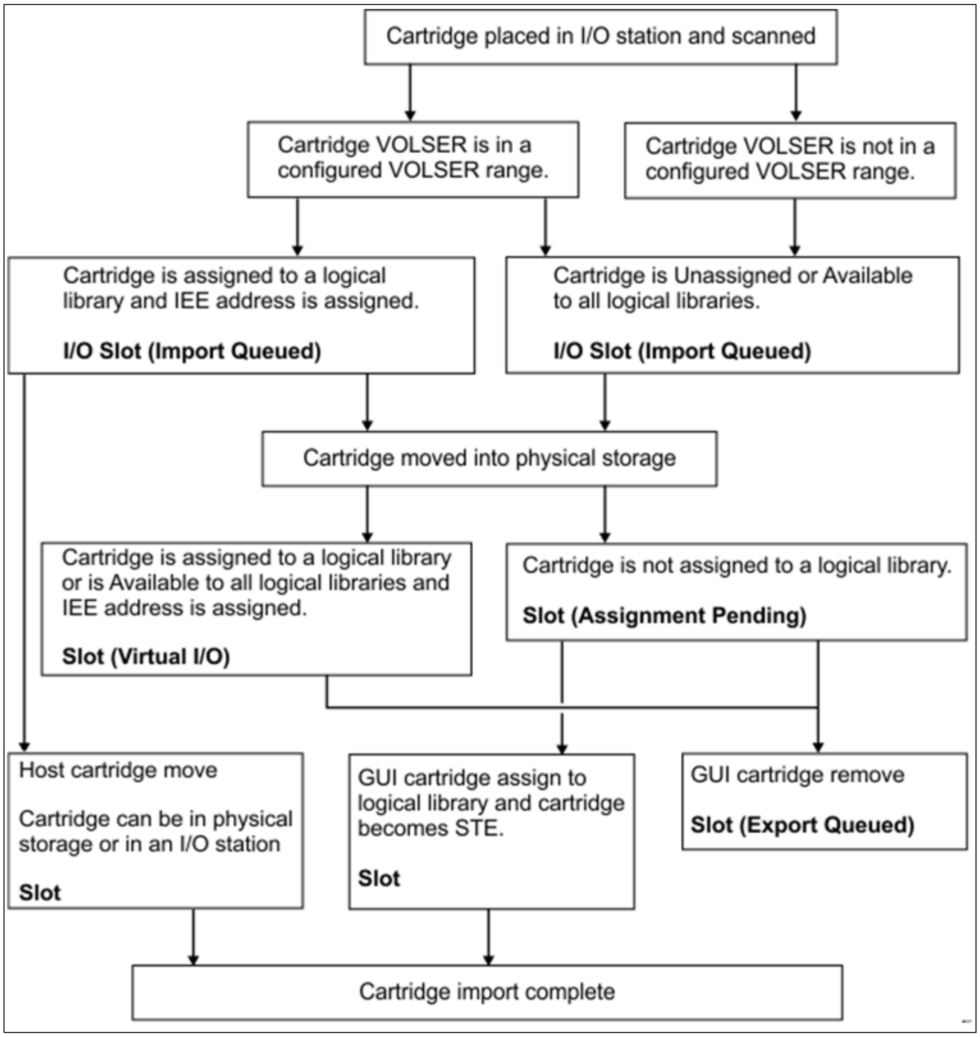

Figure 1-39 I/O import sequence of a cartridge

TS4500 integrated management console

The IMC is a built-in platform for tools that are used to manage the TS4500 tape library.

The IMC, which includes an LCD panel and a keyboard with a touchpad or trackpoint, can be mounted on either end of your TS4500 tape library. A library controller card (LCC) and a power source are required within that end frame or within the adjacent frame. Alternatively, FC 2737 allows for the IMC to be mounted on a non-powered end frame; that is, more than one frame away from a powered frame. This feature can be installed during the installation of expansion frames in an initial library installation, or later when expansion frames are added to a library.

The IMC comes preinstalled with a system console application, which is a set of software tools that are used for local service and remote support of the attached TS4500 tape library. The system console application enables the IMC to provide service console capabilities, such as broadband Call Home.

The TS4500 management GUI runs on a web browser in kiosk mode on the IMC. *Kiosk mode* means that the menu bar, address bar, and stop and reload buttons of the browser are disabled. In addition, it is not possible to use bookmarks or multiple browser windows.

Figure 1-40 shows the IMC.

Figure 1-40   Integrated management console

For more information about the IMC, see 4.1, “Integrated management console” on page 186.

IMC power distribution unit

The TS4500 IMC power distribution unit (IMC PDU) provides power to the card cage power supply and to the TS4500 IMC.

The IMC PDU, as shown in Figure 1-41, is typically installed in the end frame of the library with the IMC or last frame with power.


Figure 1-41   IMC PDU

A new Model Sx5 feature (FC 2737) provides a separate IMC power source, which enables the IMC to be installed on an Sx5 frame where the power cable is not long enough to reach a Dx5 or Lx5 frame. The feature includes instructions for moving the IMC PDU from the Lx5 frame to the Sx5 frame with the IMC, and for attaching a separately ordered power cord to a client outlet at that Sx5 frame. The feature also includes two lengthy Ethernet cables to connect the IMC to the Lx5 frame LCC at any frame position.

Power cable hole

It is possible to route power cables through the top of a frame in the TS4500 tape library. This routing method is an alternative to routing the power cables through the bulkhead at the bottom rear of the TS4500 tape library. This routing method might be necessary if the equipment is installed on a non-raised (solid) floor or if a top rack is installed.

Fibre Channel cable hole

It is possible to route Fibre Channel cables from servers through the top of a frame in the TS4500 tape library. This routing method is an alternative to routing the Fibre Channel cables through the bulkhead at the bottom rear of the TS4500 tape library. This routing method might be necessary if the equipment is installed on a non-raised (solid) floor.

TS4500 frame control assembly

The frame control assembly (FCA) is standard on all base frames (Lx5) and optional on any Dx5 expansion frames. The FCA includes one library controller card (LCC), up to two library frame interconnect (LFI) cards, and two power supplies, both of which can provide power to the library and all drives in a frame. All of these components are connected by using two back plane cards (BPCs).

Figure 1-42 shows the FCA in the TS4500.

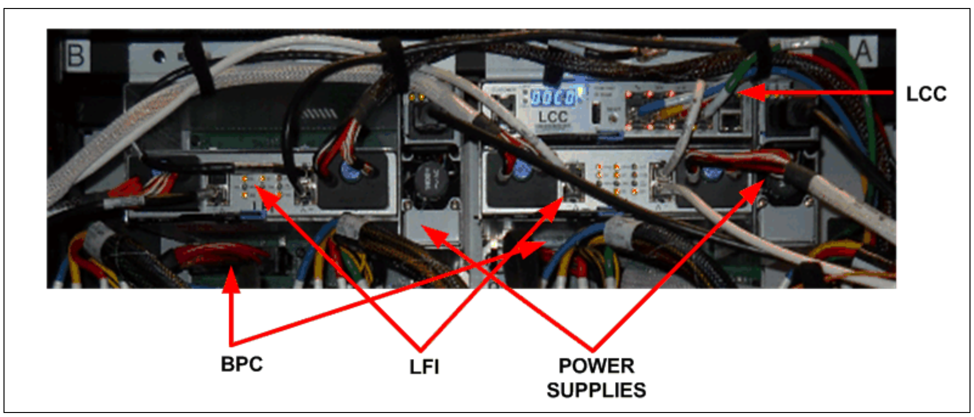

Figure 1-42 Frame control assembly

The library control card

The LCC is the management node card for the TS4500 tape library. This card is used for the Ethernet connection to all components in the library, including the user interface and IMC (see Figure 1-43).

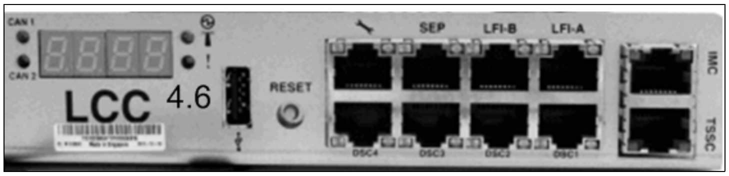

Figure 1-43 Library control card

The LCC has redundancy capability. In multiple drive frame configurations, the library negotiates the primary LCC, which controls the management functions.

During the power-up process, a selection process occurs to select the primary LCC. The selection process considers the LCC with the best database capabilities and the best hardware capabilities. The hardware capabilities relate to I/O station control and frame types.

Each LCC knows how many LCCs are in the system. If only one LCC is in the system, this LCC becomes the primary LCC. If any primary LCC fails, any LCC card can resume operations as the primary LCC.

Note: Every Lx5 frame and a Dx5 frame with drives installed have an LCC installed. Sx5 frames or Dx5 frames without drives do not have an LCC installed.

Library frame interconnect cards

The LFI cards provide communication between the frames of the TS4500 tape library through cable connections. Each Lxx frame includes one LFI card. A second LFI is required if expansion frames are added. When you face the back of the frame, the LFI on the right is LFI-A and the LFI on the left is LFI-B. Both cards are identical, but connections differ depending on whether the card is LFI-A or LFI-B.

Note: From firmware version R1.3 and higher, two LFI cards must be installed in every frame.

Figure 1-44 shows the LFI card.

Figure 1-44 Library frame interconnect

Note: Do not use tools to tighten the screws on the LFI card connectors. The use of tools can result in cables not being correctly plugged or damage to the connectors.

Two versions of LFI cards are available: the new version is labeled R1.5 or higher. The LED behavior is different since the 1.5 versions. Consider the following points:

- Previous to V1.5, the SRC LED was always on. From version R1.5, the SRC LED is only on if the LFI is creating 40 V. Therefore, it is only on in select powered frames (up to three frames in a system) and is never on in an S-frame.
- The CAN LED on the previous version of LFI is replaced with an L-frame indicator, which is labeled LFRM. The LFRM LED is on, regardless of whether the machine is turned on, if AC power is supplied to the L-Frame. This LED is on in the L-frame only.
- On previous version LFI, the Ethernet LEDs were always on. When the Ethernet cable was plugged in, the LEDs changed color. On the R1.5 version, the Ethernet LEDs are off until plugged in. After it is plugged in, the right LED turns on solid, the link LED and the left LED flash, and the activity LED flashes during data transfer.

Power supplies

For models L25, D25, L55, and D55, the library offers the TS4500 frame control assembly power structure, which combines drive power, library power, and AC power cord capabilities. Each of the Lx5 and Dx5 frames comes with two power supplies. Model Sx5 frames do not have their own power supply. Instead, they share power with an Lx5 or Dx5 frame, as shown in Figure 1-45.

Figure 1-45 Power supply

Back plane cards

Each of the Lx5 and Dx5 frames comes with two BPCs. With a single frame configuration (or a configuration with no FCA control system), redundant power is provided for each BPC. If one BPC fails, the other BPC in the frame picks up the load for the failed BPC.

Patch panel

The patch panel that houses the cable connections to connect hosts to the tape drive is in the rear of the frame above the drives, as shown in Figure 1-46.

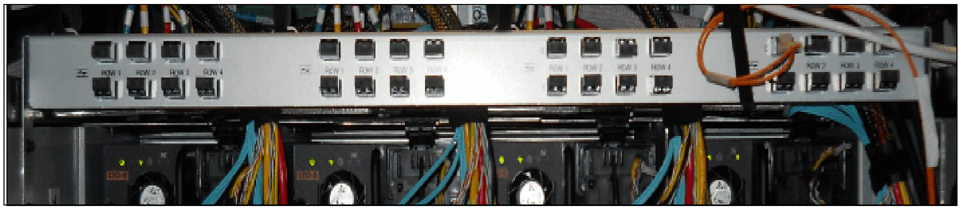

Figure 1-46 Fiber patch panel

Display panel

The display panel on the base frame of the TS4500 tape library houses the library power and pause buttons. It displays library and I/O status indicators, as shown in Figure 1-47.

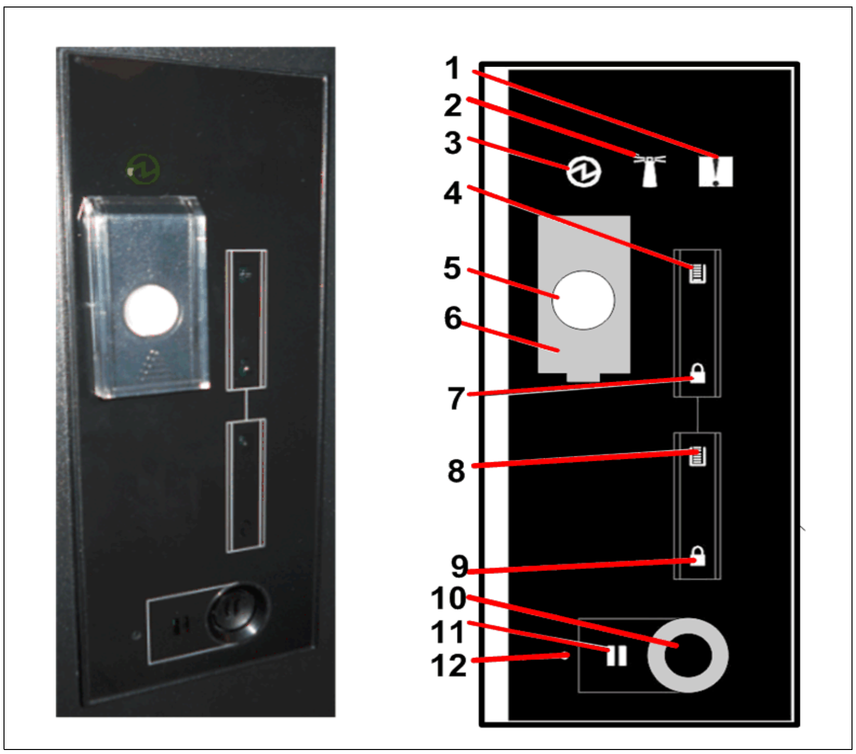

Figure 1-47 Display panel

The indicators and buttons that are available on the display panel are listed in Table 1-15.

Table 1-15 Display panel component descriptions

Number	Component
1	Warning indicator (yellow)
2	Library beacon (blue)
3	Power indicator (green)
4	Upper I/O fullness indicator (green)
5	Power switch
6	Power switch cover
7	Upper I/O lock indicator (green)
8	Lower I/O fullness indicator (green)
9	Lower I/O lock indicator (green)
10	Pause button
11	Library pause indicator (yellow)
12	Access recovery button

Expansion frames with optional installed I/O stations also feature a display panel. However, the panel on expansion frames does not include the library power button, as shown in Figure 1-48.

Figure 1-48 Expansion frame with I/O station

Pause button

You pause the library by pressing the pause button on the display panel (number 10 in Figure 1-47 on page 53) on the base frame of your TS4500 tape library. This pause button causes the library to move the cartridge accessor to the base frame and pause operations for 30 seconds or while a frame door is open. The pause light next to the pause button stops flashing, and it illuminates solid to indicate that the library is paused (number 11 in Figure 1-47 on page 53).

Power button

The power button procedures are described in this section.

Power on

Complete the following steps to power on the TS4500 tape library to begin the library initialization sequence and bring the library to a ready state:

1. From the display panel on the base frame of the library, slide the plastic cover that protects the power button up to make available the button (number 6 in Figure 1-47 on page 53).
2. Press the power button (number 5 in Figure 1-47 on page 53) and wait to see that the green power indicator turns on and stays on.

The power button runs a power-on initialization sequence for approximately 2 minutes. During that time, the menus on the TS4500 management GUI are not available for use. After the power-on initialization sequence completes, the library performs an inventory of the tape cartridges.

Power off

Complete the following steps to power off the TS4500 tape library after normal operation, but *not* during an emergency:

1. Ensure that the host application removed cartridges from all drives and that the library is varied offline from the host (if the host is attached).
2. Pause the library by pressing the pause button on the display panel (number 10 in Figure 1-47 on page 53) on the base frame of your TS4500 tape library. This pause button causes the library to move the cartridge accessor to the base frame and pause operations for 30 seconds or while a frame door is open. The pause light next to the pause button stops flashing, and it illuminates solid to indicate that the library is paused (number 11 in Figure 1-47 on page 53).

   The health status pod on the System page of the management GUI also shows when the library is paused.
3. From the display panel on the base frame of the library, slide the plastic cover that protects the power button up to make available the button (number 6 in Figure 1-47 on page 53).
4. Press the power button (number 5 in Figure 1-47 on page 53) and wait until the green power indicator light turns off.

Note: Always pause the library before you power it off. If you power off the library before you pause the library, the library might take longer to come online after the next power-on.

Access Recovery

The Access Recover button can be used to unlock the administrator password and can only be done from the front panel.

When you press the Access Recovery button, you have 15 minutes to log in to the library and reset the administrator password with a temporary password. You must enter a new administrator password when you log out and back in again.

Press the Access Recovery button (see Figure 1-47 on page 53) then, log in to the GUI by using `admin` for the user name and password.

The Library requests to change the temporal password before continuing.

Media verification

The TS4500 supports media verification, which is fully controlled by the TS4500 at a logical library level. Media verification is transparent to any software application that controls the library.

This feature can be used to verify any media for the ability to read all of the data that is on the selected media within a logical library. Media verification is set up to run automatically on a set date.

You can set a recurring schedule for verification by days, months, or years. After you set up media verification, you can verify a single medium before the set date, if required. For more information, see “Media verification” on page 258.

Web camera

After the mounting hardware feature is installed by an IBM service support representative (SSR), you can mount a web camera in any frame of a TS4500 tape library. Specific clearances are necessary to accommodate the mounting hardware and the camera.

The camera is client-supplied and it can be installed in any frame to visually monitor the location of the library’s robotics. FC 1530, Web Camera Mounting Hardware, is required to install a web camera. The maximum camera dimensions can vary. However, no part of the camera can extend more than 60 mm (2.36 in.) from the center of the tripod mount to ensure sufficient clearance so that the accessor movement is not obstructed.

The camera mount (FC 1530) can be mounted on the front or rear frame support.

Fire suppression

Each Dx5, Lx5, Sx5, and Sx4 frame in the TS4500 tape library has an area on the top where an opening can be cut to allow the entrance of pipes, conduits, or other parts to use for fire suppression. The installation of a fire-suppression system is the responsibility of the client.

The TS4500 tape library is constructed to allow the third-party installation of fire-suppression equipment. When you are deciding whether to implement fire-suppression equipment, see the local and national standards and regulations for more information.

> **Important:** For advice about selecting a fire-suppression system that provides the correct level of coverage and protection, consult your insurance underwriter and local fire marshal (or local building inspector). IBM designs and manufactures equipment to internal and external standards that require certain environments for reliable operation. Because IBM does not test any equipment for compatibility with fire-suppression systems, IBM does not make compatibility claims of any kind nor does IBM provide suggestions about fire-suppression systems.

1.4 Feature codes for the TS4500

You can order FCs to install components in library frames or to enhance the capacity or capabilities of the library.

A client-setup unit (CSU) is a feature that you, as the client, can install when the feature is ordered as a field upgrade. If you choose not to install a CSU, IBM can install it for an extra charge. However, detailed installation instructions are included when you order and receive these features. The available FCs are listed in Table 1-16.

Table 1-16 Feature codes for the frame models of the TS4500 tape library

Feature code	Model	CSU	Description
0983	Lx5, Dx5, 60F, TR1, TR2	No	TAA compliance
1404	L25, L55	No	TS1140 and TS1150 tape drive support This FC triggers the shipment of a JK diagnostic cartridge. When Model L25 is ordered, FC 1404 is required. When FC 9080 (Mixed Media) is ordered for the Model L55, FC 1404 is required.
1405	L25, L55	Both	LTO-6 and LTO-5 tape drive support This FC triggers the shipment of an L5 diagnostic cartridge. When model L55 is ordered, a minimum of one of FC 1405 or 1407 is required but both can be selected. When FC 9080 (Mixed Media) is ordered for model L25, a minimum of one of FC 1405 or 1407 is required but both can be selected.
1406	L25, L55	No	TS1160, TS1155, and TS1150 tape drive support This FC triggers the shipment of a JL CE cartridge. When model L25 is ordered or when FC 9080 (Mixed Media) is ordered for model L55 a minimum of one of FC 1404 or 1406 is required but any can be selected.
1407	L25, L55	No	LTO-7 and LTO-6 tape drive support This FC triggers the shipment of an L6 CE cartridge. When model L55 is ordered, a minimum of one of FC 1405 or 1407 is required but both can be selected. When FC 9080 (Mixed Media) is ordered for model L25, a minimum of one of FC 1405 or 1407 is required but both can be selected.
1408	L25, L55	No	LTO-8 and LTO-7 tape drive support This FC triggers the shipment of an L7 CE cartridge. When model L55 is ordered, a minimum of one of FC 1407 or 1408 is required but both can be selected. When FC 9080 (Mixed Media) is ordered for model L25, a minimum of one of FC 1407 or 1408 is required but both can be selected.

Feature code	Model	CSU	Description
1442	D25, D55, L25, L55	No	Dual accessors This FC triggers the shipment of the following components: ▶ Four CE/diagnostic cartridges (two cartridges of the latest generation for both LTO and 3592). ▶ Service Bay Stop-A-related parts and Service Bay Stop-B-related parts. ▶ TS4500 Accessor B, including scanner, enhanced grippers, and enhanced node cards. A maximum of one HA kit (FC 1442) per library can be ordered.
1450	D25, D55	No	TS4500 frame control assembly (FCA) This FC Includes 1x LCC and 2x power supplies.
1460	L25, L55	No	Redundant Accessor Power/Network This feature provides power and internal network redundancy for the L25/L55 frame. This is a prerequisite to FC 9002.
1530	D25, D55, L25, L55, S25, S55	No	Web camera mounting hardware. Provides mounting hardware only.
1531	D25, D55, L25, L55	No	First quad drive mounting kit This feature provides signal and power cabling for a column of 4 adjacent drives. Fiber cables can be ordered by using one of the following FCs: ▶ FC 1536, MMF Quad Drive-to-Patch Panel cables ▶ FC 1537, SMF Quad Drive-to-Patch Panel cables ▶ FC 9713, No Patch Panel Cables Prerequisites: ▶ Models L25/D25: One FC 1536 or 9713 is required ▶ Models L55/D55: One FC 1536 or 1537 or 9713 is required. ▶ Models D25/D55: FC 1450 is required
1532	D25, D55, L25, L55	No	Second quad drive mounting kit This feature provides signal and power cabling for a column of 4 adjacent drives. Fiber cables can be ordered by using one of the following FCs: ▶ FC 1536, MMF Quad Drive-to-Patch Panel cables ▶ FC 1537, SMF Quad Drive-to-Patch Panel cables ▶ FC 9713, No Patch Panel Cables This feature is not supported in the leftmost frame of a dual accessor library. Prerequisites: ▶ Models L25/D25: One FC 1536 or 9713 is required ▶ Models L55/D55: One FC 1536 or 1537 or 9713 is required ▶ FC 1521 or 1531 is required

Feature code	Model	CSU	Description
1533	D25, D55, L25, L55	No	Third quad drive mounting kit This feature provides signal and power cabling for a column of 4 adjacent drives. Fiber cables can be ordered by using one of the following FCs: ► FC 1536, MMF Quad Drive-to-Patch Panel cables ► FC 1537, SMF Quad Drive-to-Patch Panel cables ► FC 9713, No Patch Panel Cables This feature is not supported in the leftmost or rightmost frame of a dual accessor library. Prerequisite: ► Models L25/D25: One FC 1536 or 9713 is required ► Models L55/D55: One FC 1536 or 1537 or 9713 is required. ► FC 1522 or 1532 is required
1534	D25, D55, L25, L55	Both	Fourth quad drive mounting kit This feature provides signal and power cabling for a column of 4 adjacent drives. Fiber cables can be ordered by using one of the following FCs: ► FC 1536, MMF Quad Drive-to-Patch Panel cables ► FC 1537, SMF Quad Drive-to-Patch Panel cables ► FC 9713, No Patch Panel Cables This feature is not supported in the leftmost frame of a single accessor library. This feature is not supported in the leftmost or rightmost frame of a dual accessor library. Prerequisites: ► Models L25/D25: One FC 1536 or 9713 is required ► Models L55/D55: One FC 1536 or 1537 or 9713 is required ► FC 1523 or 1533 is required
1536	D25, D55, L25, L55	Both	MMF Quad Drive-to-Patch Panel Cables This feature includes all fiber cables that are required for a column of four drives, 8x LC-LC Multimode Fibre (MMF) Drive-to-Patch Panel cables. Prerequisite: One FC 1536 or 1537 or 9713 per each FC 1531, 1532, 1533, and 1534.
1537	D55, L55	Both	SMF Quad Drive-to-Patch Panel Cables This feature includes all fiber cables that are required for a column of four drives, 8x LC-LC Single Mode Fibre (SMF) Drive-to-Patch Panel cables. Prerequisite: One FC 1536 or 1537 or 9713 per each FC 1531, 1532, 1533, and 1534.
1604	L25, L55	Yes	Transparent LTO encryption Provides license keys to enable transparent LTO encryption.
1628	L55	Yes	Additional LTO Cartridge Magazines Provides magazine and cover for transport or storage of up to 18 cartridges.

Feature code	Model	CSU	Description
1629	L25	Yes	Additional 3592 Cartridge Magazine Provides magazine and cover for transport or storage of up to 16 cartridges.
1643	L25, L55	Yes	Intermediate CoD Provides a license key to increase storage from entry capacity to intermediate capacity.
1644	L25, L55	Yes	Base CoD Provides a license key to increase storage from intermediate capacity to base capacity. This feature code is a prerequisite for FC 9002 or Model Lx5 HD CoD FC 164x. Prerequisite: FC 1643.
1645	S25	Yes	High-density CoD
1646	S55	Yes	High-density CoD
1647	L25	Yes	High-density CoD Prerequisite: FC 1644.
1648	L55	Yes	High-density CoD Prerequisite: FC 1644.
1649	D25	Yes	High-density CoD
1650	D55	Yes	High-density CoD
1652	D25, D55	No	Two additional I/O stations. This feature adds 36 (LTO) or 32 (3592) I/O slots.
1663	D25, D55, L25, L55	No	Drive removal One 3592 tape drive installation feature or one 3588 tape drive installation feature should be removed from the library frame when a tape drive is removed.
1682	L25, L55	Yes	Path failover
1742	S24, S54	No	TS4500 Control upgrade This FC includes 1x top panel and 2x Ethernet cables, card cages, ASCs, BPCs, and LFIs.
1750	TR1, TR2	No	Left-side and right-side end covers This feature is required only for the first top rack that is ordered (if multiple top racks are installed on adjacent frames).

Feature code	Model	CSU	Description
1751	TR1	No	Power distribution unit Provides one PDU to be mounted in the top rack. The first PDU in a top rack does not consume any of the 10U of rack space. The second PDU in a top rack consumes 1U of rack space. Consider the following points: ► A maximum of two FC 1751s can be ordered. ► FC 1751 cannot be installed in the same rack as an FC 1752 PDU. Each PDU has 10 C13 outlets that are divided into two groups of five outlets. Each group is protected by a 15 amp circuit breaker. In addition, a master 25 amp circuit breaker is available for the entire PDU. All of the circuit breakers are rated at 5 kAIC. Each PDU has a control input that can be used to turn all of the outlets on or off by an external switch or relay contact (not provided). The control input is defaulted to ON by a jumper plug that is included with the PDU. Prerequisite: For each FC 1751 ordered, one power cord (FC 9954 - 9959 or 9966) must also be ordered. Limitation: Mutually exclusive with 1752
1752	TR1, TR2	Yes	Enhanced PDU Provides one PDU with 3-phase (wye) or single phase input power to be mounted in top rack (max of two can be ordered). The first PDU in a top rack does not use any of the 10U of rack space. The second PDU in a top rack consumes 1U of rack space. Each PDU has nine outlets that are divided into three groups. Within each group, two C-13s and one C-19 outlets are used. Each group is protected by 20 amp circuit breakers that are rated at 10 kAIC. If required, loads that are installed in the TR1 with C14 inlets can be connected to the PDU C19 outlets by using C20 - C13 power cord FC 9949. Prerequisites: ► For single phase input power, for each PDU ordered must order one power cord FC 9954 - 9958. ► For 3-phase (wye) input power, for each PDU ordered must order one power cord FC 9948. Limitations: Mutually exclusive with 1751

Feature code	Model	CSU	Description
1753	TR2	No	Olympus Rack Configuration This feature specifies the configuration of the TR2 rack to be in conformance with the Open Compute Project Olympus Rack Specification. This provides 5U of 19 inch rack space with three sets of EIA rails with a spacing from the front face of the front EIA rail to the rear face of the middle EIA rail being 745 mm (29.3 inches) and the spacing from the front face of the front EIA rail to the rear face of the rear EIA rail to be 1020 mm (41 inches). It includes a top cover that allows top access cables to be routed into the TR2. This also includes side air blocks that prevent airflow between the EIA rails and the side covers per the Open Compute Olympus Rack Specification.
1754	TR2	No	Standard 19 inch Rack Configuration This feature specifies the configuration of the TR2 rack to be in conformance with industry standard 19-inch racks. This provides 5U of 19 inch rack space with two sets of EIA rails with a spacing from the front face of the front EIA rail to the rear face of the rear EIA rail being 719 mm (28.3 inches). It includes a top cover that allows top access cables to be routed into the TR2.
1755	TR2	Yes	Front Door This feature provides the front door for the TR2
1756	TR2	Yes	Rear Door This feature provides the rear door for the TR2
1802	L25, L55	No	Non-HA 1 - 2 frame X-Track Cable
1806	L25, L55	No	Non-HA 3 - 6 frame X-Track Cable
1814	L25, L55	No	Non-HA 7 - 14 frame X-Track Cable Non-HA 7 to 14 Frame Track Cable. For more information, see 3584 X-Track Cable Features/Specify codes. Limitation: Not supported by FC 2071 or FC 9071.
1818	L25, L55	No	Non-HA 9 - 18 frame X-Track Cable Non-HA 9 to 18 Frame Track Cable. For more information, see 3584 X-Track Cable Features/Specify codes. Limitation: Not supported by FC 2071 or FC 9071.
1909	L25, L55, D25, D55	No	Single power source bifurcated cable This feature allows attachment of a powered frame to a single outlet while it maintains drive redundant power.

Feature code	Model	CSU	Description
1951	D25, D55, L25, L55, S25, S55	No	Power distribution units Provides two PCAs, mounting hardware including rack brackets, and internal power cables. Supports any combination of up to three pairs of power cords for FC 4875 on the same frame and FC 9989 on adjacent frames. Each PDU has 10 C13 outlets that are divided into two groups of five outlets. Each group is protected by a 15 amp circuit breaker. In addition, a master 25 amp circuit breaker is available for the entire PDU. All of the circuit breakers are rated at 5 kAIC. Each PDU has a control input that can be used to turn all of the outlets on or off by an external switch or relay contact (not provided). The control input is defaulted to ON by a jumper plug that is included with the PDU. Corequisite: One power cord feature (FC 9954 - 9959 or 9966) Limitation: FC 1951 cannot be installed in the same rack as an FC 1952 PDU.
1952	D25, D55, L25, L55, S25, S55	No	Enhanced PDUs Provides two PDUs with 3-phase (wye) or single phase input power, mounting hardware including rack brackets, and internal power cables to provide power to the local frame control assembly. These power cord features might require different customer facility outlets than frames without FC 1952. Can be used for adjacent frame AC power cord aggregation, or for unique facility AC socket requirements. Supports any combination of up to three pairs of power cords for FC 9989 on adjacent frames. Each PDU has nine outlets that are divided into three groups. Within each group, two C-13s and one C-19 outlets are used. Each group is protected by 20 amp circuit breakers that are rated at 10 kAIC. If required, loads that are installed in the TR1 with C14 inlets can be connected to the PDU C19 outlets by using C20 - C13 power cord FC 9949. Prerequisites: ► For single phase input power, must order one power cord FC 9954 - 9958. ► For 3-phase (wye) input power, must order one power cord FC 9948. Limitations: ► Mutually exclusive with legacy PDU FC 1951, Fibre Channel switch mounting FC 4879, and Ethernet switch FC 2704. ► Mutually exclusive with FC 4879 or 2704.
2002	L25, L55	No	1 - 2 Frame Flex Track Cable A
2006	L25, L55	No	3 - 6 Frame Flex Track Cable A
2014	L25, L55	No	7 - 14 Frame Flex Track Cable A
2018	L25, L55	No	9 - 18 Frame Flex Track Cable A

Feature code	Model	CSU	Description
2071	L25, L55	No	Flex track guide This FC includes AXY, AXY plate, ASC to Flex track cable and new end stops.
2072	D25, D55, S25, S55, S24, S54	No	Flex track guide for expansion frame
2309	D25, D55, L25, L55	No	TS4500 Accessor Refresh This FC provides complete accessor assembly with HD grippers to Lxx models or models with FC 1442. Prerequisite: FC 2071, FC 2072, or FC 9071, FC 9072.
2402	L25, L55	No	1 - 2 Frame Flex Track Cable B
2406	L25, L55	No	3 - 6 Frame Flex Track Cable B
2414	L25, L55	No	7 - 14 Frame Flex Track Cable B
2418	L25, L55	No	9 - 18 Frame Flex Track Cable B
2704	D25, D55, S25, S55, S24, S54	No	Console expansion 26-port Ethernet switch, rack mount Provides a 26-port Ethernet switch and attachment cable for connection to an IMC as a centralized system console. Up to 24 extra connections of FC 2715 are provided by this feature. This feature is supported on Lx5 models to allow other products to share the IMC.
2715	L25, L55	No	TSSC attachment cable FC 2715 is a cable to attach a unit to the Ethernet switch that is provided by the TSSC. Note: ▶ A maximum of 43 of FC 2715 can be included in a single TSSC facility. ▶ FC 2704 is supported on Lx5 models if you are connecting the frame to a TSSC.
2735	L25, L55	No	USB Optical Drive Required for all new plant orders. This feature provides a USB optical drive for use with the integrated management console (IMC). The optical drive is required for IMC code and service IBM Documentation updates.
2737	S25, S55	No	IMC Separate Power Source. This feature allows longer distances between the IMC end panel and an FCA by supplying 2 Ethernet cables that are longer than normal. For both plant and field, instructions are also provided for moving the IMC PDU from another frame into this frame. Prerequisite: One of the FC 9970-9985, 9989 power cords.

Feature code	Model	CSU	Description
4879	L25, D25, TR1, TR2	No	TS7700 BE Switch Mounting Hardware This feature provides 2x Fibre Channel switch mounting hardware, including rack brackets and power cords. Prerequisite: FC 4880
4880	L25, D25, TR1, TR2	No	TS7700 BE 16 Gb Switch Provides 1x Fibre Channel switch (16 Gb). Limitation: A maximum of 2x FC 4880 per model x25 can be used.
4892	D25, D55, L25, L55, S25, S55	No	Rear Door Sensor This feature provides rear door sensor components for notifying of rear door open/close events.
6013	D25, D55, L25, L55	Yes	LC to LC Fibre Channel cable - 13 m (43 ft.)
6025	D25, D55, L25, L55	Yes	LC to LC Fibre Channel cable - 25 m (82 ft.)
6061	D25, D55, L25, L55	Yes	LC to LC Fibre Channel cable - 61 m (200 ft.)
8750	D55, L55	No	LTO cleaning cartridge
8802	D25, L25	No	3592 cleaning cartridge
9001	D25, D55, L25, L55	No	Driveless frame
9002	L25, L55	No	First expansion frame attachment This feature is used as a part of the process to add any Dxx or Sxx frame. Prerequisites: ► FC 1460 and FC 1644 ► FC 9001 for Lxx models
9003	L25, L55	No	Additional expansion frame attachment This feature is used as a part of the process to add any Dxx or Sxx frame.
9040	L25, L55	No	High availability library Prerequisites: One Flex-track A and B cable of the same length Flex-Track A cable: FC 2002, 2006, 2014, or 2018 Flex-Track B Cable: FC 2402, 2406, 2414, or 2418
9071	L25, L55	No	Flex Track Guide Base This feature includes AXY, AXY plate, ASC to Flex track cable and new end stops. Prerequisite: FC 2002 or FC 2006 or FC 2014 or FC 2018.
9072	D25, D55, S25, S55, S54, D54	No	Flex Track Guide for expansion frame

Feature code	Model	CSU	Description
9073	D25, D55, L25, L55, S25, S55	No	3-Phase PDU Capable This feature indicates that the frame can support Enhanced PDUs and their associated power cords. FC 9073 is a prerequisite for ordering FC 1952.
9080	L25, L55	No	Mixed Media Library. Prerequisite: Library contains a mix of LTO and 3592 frames. For Model L25, FC 1628 (LTO Magazine) and a minimum of one of FC 1405 or FC 1407 (LTO CE cartridge).For Model L55, FC 1629 (3592 Magazine) and FC 1404 (3592 CE cartridge). Limitation: Mixed media is not supported for Lx2/Lx3.
9210	L25, L55	No	Attached to HP-UX system
9211	L25, L55	No	Attached to Solaris system
9212	L25, L55	No	Attached to Microsoft Windows system
9213	L25, L55	No	Attached to another system that is not an IBM system
9215	L25, L55	No	Attached to Linux system
9217	L25, L55	No	Attached to z/OS, TS7700
9218	L25, L55	No	Attached to High Performance Storage System (HPSS)
9400	L25, L55	No	Attached to IBM i5/OS or IBM OS/400® system
9600	L25, L55	No	Attached to IBM AIX® system
9607	D55, L55	No	3588 F7C plant installation
9609	D55, L55	No	3588 F8C plant installation
9611	D55, L55	No	3588 F8S plant installation
9613	D55, L55	No	3588 F9C plant installation
9615	D55, L55	No	3588 F9S plant installation
9617	D55, L55	No	3588 S9Cplant installation
9677	D25, L25	No	Plant Installation 3592 in a 3584 frame This code notifies the plant to factory installation a new 3592 tape drive into a new 3584 tape library frame that is coming from the plant.
9689	D25, L25	Yes	Field Installation 3592 in a 3584 frame This feature notifies the plant to ship one 3592 tape drive, which is field that is installed in a 3584 tape library frame.
9690	D25, D55, L25, L55	No	3588/3592 drive. Field install. FC 9690 is a counting feature code. Prerequisite: FC 152x (drive mounting kit). Each FC 152x supports up to four of FC 969x.

Feature code	Model	CSU	Description
9692	L25, L55	No	3592 E07/EH7 tape drive. Plant install. FC 9692 is a manufacturing routing code. Prerequisite: FC 152x (drive mounting kit).
9695	L55, D55	No	3588 F5A/F5C tape drive. Plant install. FC 9695 is a manufacturing routing code. Prerequisite: FC 152x (drive mounting kit).
9697	L55, D55	No	3588 F6A/F6C tape drive. Plant install. FC 9697 is a manufacturing routing code. Prerequisite: FC 152x (drive mounting kit).
9699	L25, D25	No	3592 EH8 tape drive - plant install FC 9699 is a manufacturing routing code. Prerequisite: FC 152x (drive mounting kit).
9700	D25, D55, L25, L55	No	No host attach cables (from plant)
9704	L55, D25	No	3592 55E tape drive. Plant install. FC 9704 is a manufacturing routing code. Prerequisite: FC 152x (drive mounting kit).
9705	L55, D25	No	3592 55F tape drive. Plant install. FC 9705 is a manufacturing routing code. Prerequisite: FC 152x (drive mounting kit).
9706	D25, L25	No	3592 60F tape drive - Plant Install in 3584
9707	D25, L25	No	3592 60E tape drive - Plant Install in 3584
9708	D25, L25	No	3592 60S tape drive - Plant Install in 3584
9713	D25, D55, L25, L55	Both	Direct Fiber - No Patch Panel Cables This feature is specified whether you do not want the factory to ship any Fibre Channel cable FC 1536 “MMF Quad Drive-to-Patch Panel cables” or FC 1537 “SMF Quad Drive-to-Patch Panel cables”.
9726	D25, L25	Both	TS1160 tape drive - Field Install Drive in 3584 This feature field installs one 3592 Model 60E, 60F, or 60S tape drive into an installed 3584 Model D25 or L25. This feature is also required to field merge a 3592 Model 60E, 60F or 60S tape drive into a 3584 frame coming from the plant.
9735	L25, L55	No	IMC broadband Call Home This feature originally indicated that the customer allows broadband Call Home; however, the feature code is no longer used.
9948	D25, D55, L25, L55, S25, S55, TR1, TR2	Both	3-Phase Power Cord This feature provides a 4.3-meter (14-foot) long power cord with an IEC 309 3P+N+G 32A plug, which is rated for 230 Vac, 24 Amps. This power cord supports 3-phase (wye) power. To be used with Enhanced PDU FC 1752 or FC 1952.

Feature code	Model	CSU	Description
9949	D25, D55, L25, L55, S25, S55, TR1, TR2	Both	C20-C13 Power Cord for spare outlets in the Enhanced PDU This feature provides a 2.8 meters long C20-C13 Power Cord to be used between computer equipment and the Enhanced PDU Feature 1752 or 1952 or customer-supplied PDU with C19 outlets. C13 connects to the male C14 inlet or power cord plug and the C20 mates with the C19 receptacle on the Enhanced PDU. This power cord is intended to be used on PDUs with C19 outlets.
9954	D25, D55, L25, L55, S25, S55, TR1, TR2	No	NEMA L6-30 Power Cord FC 9954 is used with PDU FCs 1751 or 1951. FC 9954 provides 1x power cord for FC 1751 and 2x power cords for FC 195x.
9955	D25, D55, L25, L55, S25, S55, TR1, TR2	No	RS 3750DP power cord FC 9955 is used with PDU FCs 1751 or 1951. FC 9955 provides 1x power cord for FC 1751 and 2x power cords for FC 195x.
9956	D25, D55, L25, L55, S25, S55, TR1, TR2	No	IEC 309 power cord FC 9956 is used with PDU FCs 1751 or 1951. FC 9956 provides 1x power cord for FC 1751 and 2x power cords for FC 195x.
9957	D25, D55, L25, L55, S25, S55, TR1, TR2	No	4.3 m power cord (Australia/NZ) FC 9957 is used with PDU FCs 1751 or 1951. FC 9957 provides 1x power cord for FC 1751 and 2x power cords for FC 195x.
9958	D25, D55, L25, L55, S25, S55, TR1, TR2	No	4.3 m power cord (Korean) FC 9958 is used with PDU FCs 1751 or 1951. FC 9958 provides 1x power cord for FC 1751 and 2x power cords for FC 195x.
9959	D25, D55, L25, L55, S25, S55	No	Unterminated power cord FC 9959 is used with PDU FCs 1751 or 1951. FC 9959 provides 1x power cord for FC 1751 and 2x power cords for FC 195x.
9966	D25, D55, L25, L55, S25, S55	No	Unterminated power cords (China CCC cert) FC 9966 is used with PDU FCs 1751 or 1951. FC 9966 provides 1x power cord for FC 1751 and 2x power cords for FC 195x.

Feature code	Model	CSU	Description
9970	D25, D55, L25, L55, S25, S55	No	Dual 4.3 m power cord (international, 250 VAC 16A single phase, watertight connector, IEC-309) FC 9970 includes model-dependent contents. For countries other than the US and Canada, the connector is rated at 16A (plug type Hubbell HBL316P6W or equivalent) for connection to Hubbell type HBL316R6W or equivalent receptacles. In the US and Canada, the connector is rated at 20 A (plug type Hubbell HBL320P6W) for connection to Hubbell type HBL320R6W or equivalent receptacles. This power cord is the default power cord for all countries other than those countries that are specified for FC 9972 and FCs 9976 - 9983. This power cord can be used for the following countries: ► Argentina ► Australia ► Brazil ► Canada ► China ► Japan ► Korea ► New Zealand ► Philippines ► South Africa ► Taiwan Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC 2737
9972	D25, D55, L25, L55	No	Dual 4.3 m power cords (250 VAC 15, non-watertight twist lock connector, Nema L6-15P mates with L6-15R). FC 9972 includes model-dependent contents. It is the default power cord for US, Canada, Japan, Korea, Philippines, and Taiwan. Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9976	D25, D55, L25, L55	No	Dual 4.3 m power cords with 10 Amp/250 Vac, non-watertight IRAM 2073 plug (Argentina). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737
9977	D25, D55, L25, L55	No	Dual 4.3 m power cords with 15 Amp/250 Vac with earth pin InMetro NBR 14136 plug (Brazil). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9978	D25, D55, L25, L55	No	Dual 4.3 m power cords with 10 Amp/250 Vac AZ/NZS 3112/2000 plug (Australia and New Zealand). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.

Feature code	Model	CSU	Description
9979	D25, D55, L25, L55	No	Dual 4.3 m power cords with 15 Amp/250 Vac JIS C8303, C8306 plug (Japan). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9980	D25, D55, L25, L55	No	Dual 4.3 m power cords with 10 Amp/250 Vac, GB 2099.1, 1002 plug (China). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9981	D25, D55, L25, L55	No	Dual 4.3 m power cords with 15 Amp/250 Vac with earth pin KS C8305, K60884-1 plug (Korea). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9982	D25, D55, L25, L55	No	Dual 4.3 m power cords with 10 Amp/250 Vac CNS 10917-3 plug (Taiwan). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9983	D25, D55, L25, L55	No	Dual 4.3 m power cords with 10 Amp/250 Vac SANS 164-1 plug (South Africa). Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9984	D25, D55, L25, L55	No	Dual 4.3 m power cords with 15 A/250V single phase power cord assemblies, NEMA L6-20P non-watertight twistlock 20A plug. FC 9984 is 2x power cords, which are UL/CSA certified for use in US and Canada. These power cords mate with client-supplied NEMA L6-20R receptacles. Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9985	D25, D55, L25, L55	No	Dual 4.3 m power cords with 15 A/250V single phase power cord assemblies with watertight 15A Russellstoll plug. FC 9985 is 2x power cords, which are UL/CSA certified for use in US and Canada. These power cords mate with client-supplied Russellstoll 3743U2 or 9R23U2W receptacles. Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.
9989	D25, D55, L25, L55	No	Dual 4.3 m power cords with 10A/250V with IEC 309 C14 plug. FC 9989 is for use with adjacent frame PDUs (FC 1951) or external (client-supplied) PDUs. Note: Only 1 power cord is shipped for an S25 and S55 to be used with FC2737.

Feature code	Model	CSU	Description
AGKQ	L25, D25, L55, D55	No	3 Meter OM3 Fiber Quad Cables (LC) Provides two space-efficient 3 m (9.8-foot) bundles (of four cables each) to connect four drives to a fiber switch in the same or one frame away, or in a TR1 top rack. To attach tape drives in the same frame or 1 frame away from a model x25 frame containing TS7700 BE switches, order one of this feature code for each quaddrive mounting kit to be connected. For xx5 models, each FC AGKQ has a prerequisite of an associated FC 152x or 153x.
AGK1	D25, D55, L25, L55	No	10-meter OM3 fiber Cable (LC) This feature provides a 10-meter (32.8-foot) 50.0/125 micrometer fiber optic cable that is terminated with LC Duplex connectors.
AGK2	D25, D55, L25, L55	No	25-meter OM3 fiber Cable (LC) This feature provides a 25-meter (82-foot) 50.0/125 micrometer fiber optic cable that is terminated with LC Duplex connectors.
AGK3	D25, D55, L25, L55	No	80-meter OM3 fiber Cable (LC) This feature provides an 80-meter (262-foot) 50.0/125 micrometer fiber optic cable that is terminated with LC Duplex connectors.

For more information about ordering media, see Ordering LTO cartridges and media supplies and Ordering 3592 cartridges and media supplies in IBM Documentation.

1.5 Host platforms and device drivers

The TS4500 tape library is supported on many operating systems. For a current list of host software versions and release levels that support the TS4500 tape library, see the IBM System Storage Interoperation Center (SSIC) web page to confirm support and to check for all required firmware and drivers.

To download the required drivers, see IBM Fix Central.

Feature codes for operating systems

The following no-charge specify feature codes indicate the server platform to which the TS4500 tape library is attached. These features are used by IBM for device driver distribution:

- FC 9210: Attached to HP-UX
- FC 9211: Attached to Sun system
- FC 9212: Attached to Windows system
- FC 9213: Attached to another system that is not an IBM system
- FC 9215: Attached to Linux system
- FC 9217 Attached to IBM System z®, TS7700
- FC 9218: Attached to High Performance Storage System (HPSS)

- FC 9400: Attached to IBM System i
- FC 9600: Attached to IBM System p

More than one platform-attached features can be chosen because the library can be attached to more than one of these platforms. Only one of each feature is added; that is, only one FC 9212 is required if two or more Windows servers use the TS4500.

Tip: The device driver is available for each operating system from IBM Fix Central.

1.6 Specifications

The TS4500 tape library is a stand-alone tape subsystem that consists of one or more frames. It can expand in a modular manner to provide large capacities. The frames join side-to-side, and they can be added on the left side or right side of the installed frames.

1.6.1 Physical specifications

The physical dimensions of the TS4500 tape library frames are listed in Table 1-17.

Table 1-17 Physical dimensions

Model	Width[a]	Depth	Height
L25, L55, D25, D55, S25, S55, S24, and S54 (on casters)	782 mm (30.8 in.)	1,212 mm (47.72 in.)	1,800 mm (70.9 in.)
L25, L55, D25, D55, S25, S55, S24, and S54 frames with top rack	782 mm (30.8 in.)	1,212 mm (47.72 in.)	2,320 mm (91.34 in.)
L25, L55, D25, D55, S25, S55, S24, and S54 frame with covers	782 mm (30.8 in.)	1,212 mm (47.72 in.)	1,800 mm (70.9 in.)
L25, L55, D25, D55, S25, S55, S24, and S54 frame without covers	725 mm (28.5 in.)	1,212 mm (47.72 in.)	1,800 mm (70.9 in.)

a. Frame width only. Extra interframe spacing of 30 mm (1.2 in.) is required.

The weights of the frames according to the number of installed drives, robotics, and tape cartridges are listed in Table 1-18.

Table 1-18 TS4500 tape library weight

Model and position	Minimum weight of frame (with no tape drives or cartridges)[a]	Maximum weight of frame (with maximum tape drives and cartridges) [b]
L25 Position 1 (leftmost frame)	407.8 kg (899 lb.)	586.5 kg (1,293 lb.)
L25 Position 2+	407.8 kg (899 lb.)	630.5 kg (1,390 lb.)
D25 Position 1 (leftmost frame)	310.7 kg (685 lb.)	500 kg (1,103 lb.)
D25 Position 2+	310.7 kg (685 lb.)	552.5 kg (1,218 lb.)
L55 Position 1 (leftmost frame)	410 kg (904 lb.)	592 kg (1,305 lb.)
L55 Position 2+	410 kg (904 lb.)	637 kg (1,404 lb.)
D55 Position 1 (leftmost frame)	316 kg (697 lb.)	503.5 kg (1,110 lb.)
D55 Position 2+	316 kg (697 lb.)	557 kg (1,228 lb.)

Model and position	Minimum weight of frame (with no tape drives or cartridges)[a]	Maximum weight of frame (with maximum tape drives and cartridges) [b]
S25 Position 1 (leftmost frame)	299 kg (660 lb.)	491.2 kg (1,083 lb.)
S25 Position 2+	299 kg (660 lb.)	540.2 kg (1,191 lb.)
S55 Position 1 (leftmost frame)	304 kg (670 lb.)	509.4 kg (1,123 lb.)
S55 Position 2+	304 kg (670 lb.)	564 kg (1,244 lb.)
S24 Position 2+	285.8 kg (630 lb.)	526.2 kg (1160 lb.)
S54 Position 2+	290.3 kg (640 lb.)	562.5 kg (1240 lb.)
TR1 Weight of top rack (empty)[c]	24.5 kg (54 lb.)	24.5 kg (54 lb.)
Weights that are listed for the Lx5 frames include the accessor, IMC, side doors, and side panels. After the initial library installation, these parts can be moved to other frames within the library string to shift part of the weight to other frames.		

a. Frames in position 1 (the leftmost frame in a library string) can have a maximum of 12 tape drives.
b. Drive frames in positions 2+ can accommodate 16 tape drives.
c. A top rack can be installed optionally on any HD frame. Side panels and PDUs are also optional. Each side panel adds 6.8 kg (15 lb.). Each PDU adds 4.5 kg (10 lb.).

When you plan for the installation, the space implications in the computer room must be considered for the possibility of adding frames in the future.

1.6.2 Floor requirements

The library must be installed on a raised or solid floor. The floor must have a smooth surface and, if raised, the floor must not have ventilation panels beneath the leveling jackscrews. If carpeted, it must be confirmed that the carpet is approved for computer-room applications. To accommodate unevenness in the floor, the leveling jackscrews can be raised or lowered to the following specifications:

- The maximum allowable variance must not exceed 7 mm (0.27 in.) per 76 mm (3 in.).
- The maximum out-of-level condition must not exceed 40 mm (1.6 in.) over the entire length and width of the library.

The floor on which the library is installed must support the following weight specifications:

- Up to 4.8 kilograms per square cm (68.6 lbs. per square inch) of point loads that are exerted by the leveling jackscrews
- Up to 211 kilograms per square meter (43.4 lbs. per square foot) of overall floor loading

The number of point loads that is exerted depends on the number of frames that makes up the library. Four point loads are on each frame (at the corners of each frame).

For more information, see the *IBM TS4500 Introduction and Planning Guide*, SC27-5990.

1.6.3 Operating environment

For more information about environmental specifications, see this IBM Documentation web page.

1.6.4 Power and cooling specifications

Power for the components of the TS4500 tape library is provided by a power supply within the Lx5 or Dx5 frame. The Lx5 power supply provides power to the components that are housed in the Lx5 frame. The Dx5 power supply provides power to the components in the Dx5 frame. The Sx5 models receive their power supply from an adjacent frame (an Lx5 or Dx5, depending on the configuration).

Models Lx5 and Dx5 of the TS4500 tape library with installed PDUs (FC 1951) can power to two adjacent frames by specifying a power cord (FC 9989) on each of the adjacent frames and plugging the cords into the spare outlets on the (FC 1951) PDUs, which are on the adjacent Lx5 or Dx5 frame.

Figure 1-49 shows the first adjacent frame cord that is plugged into the PDU receptacles and the second adjacent frame power cord that is plugged into the consecutive PDU receptacles. The adjacent frames that receive power through FC 9989 can be up to three frames away from the frame with FC 1951. By using adjacent frame power cords, clients can power to three frames from only two facility outlets.

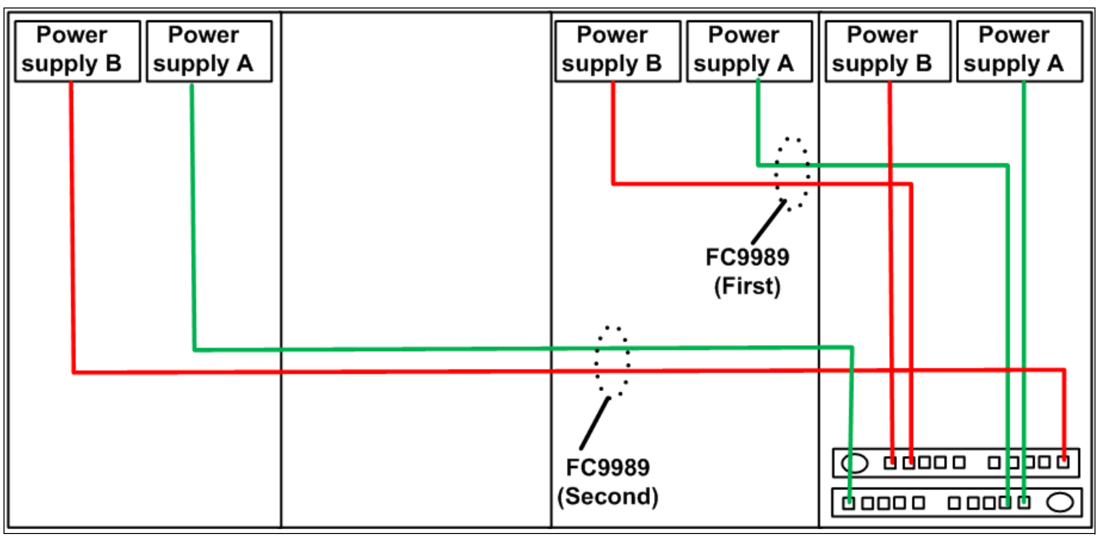

Figure 1-49 FC 9989 cable positions with FC 1951

The power requirements for the Lx5 and Dx5 frames are listed in Table 1-19.

Table 1-19 Power requirements for the TS4500 tape library

Description	**Power consumption (Watts)**			**Cooling Requirements (Btu per hr, maximum continuous)**[a] [a]
	Off [b]	**Idle**	**Maximum continuous (not peak)**	
Models, power per frame				
L25, L55[c d]	11	95	130	433
D25, D55, S25, S55[c e]	0	11	11	38
Feature Codes				
FC 1442 HA Kit w/Second Accessor	0	24	85	290
FC 1450 (TS4500 FCA) including one LCC and two AC/DC power supplies	11	38	38	130
FC 1460 (second LFI for Lx5 frame in a multiple frame network)	0	5	5	17A
FC 1521, 1523, 1524 (quad drive mounting kits)	0	3	3	10
FC 1751 PDU (for TR1 top rack)	9	9	9	31
FC 1951 PDU	17	17	17	58
Drives, power is per drive				
TS1140[f]	0	28	50	171
TS1150[f]	0	23	48	164
TS1155 55F[f]	0	23	48	164
TS1155 55E[f]	0	50	63	215
TS1160E[f]	0	38	62	211
TS1160F[f]	0	39	63	215
TS1160S[f]	0	38	62	211
TS1050[f]	0	16	32	109
TS1060[f]	0	13	36	123
TS1070[f]	0	15	31	106
TS1080[f]	0	18	40	136
TS1090[f]	0	18.5	42	143

a. To calculate the total cooling that is required by the library (in Btu per hr), multiply the total power in watts by 3.41. To convert Btu per hr to kBtu per hr, divide your result by 1000.
b. Off refers to power consumed when the library is connected to an AC power source and the library on/off switch is set to off.
c. Lx5 models and models with FC 1450 or FC1951 are equipped with dual AC power cords. The figures in the table show the total power that is consumed, including power that is consumed by redundant power supplies. Each power cord supplies approximately half of the power.
d. Lx5 frame power includes all frame loads, cartridge accessor, IMC, and display.
e. For base models that do not include FC 1450 (TS4500 FCA), the power consumption that is shown in the table figures is power that is provided by other frames. The figures include power consumption for the lighting and two LFIs.
f. Idle power is consumed when the drive has no tape cartridge loaded. Maximum continuous power is consumed when the drive is actively reading and writing to the tape. These power consumption values include the power that is required for the cooling fan at normal speed. In ambient environments that are hotter than the recommended range, the cooling fan might speed up and draw more power.

For more updated information to calculate power consumption with drives, see this IBM Documentation web page.

The environmental specifications for the TS4500, which refers to the hardware of the TS4500 tape library, and might lead to temperatures greater than allowable for the cartridges and media that are stored in the library, are listed in Table 1-20.

Table 1-20 Equipment environment specifications for the TS4500 tape library

Product operation						**Product power off**		
Dry-bulb temperature[a]		**Relative Humidity (Non-condensing)**		**Maximum Wet-bulb temp.**	**Maximum elevation**	**Dry-bulb temp.**	**Relative humidity (% RH)**	**Maximum Wet-bulb temp.**
Allowable	**Recommended**	**Allowable (% RH)**	**Recommended (% RH)**					
16 - 32°C	16 - 25°C	20 - 80%	20 - 50%	26°C	3050 m	5 - 45°C	5 - 80%	28°C

a. Derate maximum dry-bulb temperature 1°C/300 m (1.8°F/984 ft.) above 900 m (2 953 ft.).

Data processing design criteria uses a preferred controlled environment of 22° C at 45% relative humidity for the most reliable operation and performance. However, recommended operating limits extend 18 - 27° C and 40% - 60% RH or 5.5° C DP-Dew Point - 15° C DP. (ASHRAE 2016 TC9.9) Best performance is not obtained if the environment swings from limit to limit.

For frames that contain LTO tape cartridges or IBM 3592 Enterprise Tape Cartridges, see their environmental and shipping specifications for these products and adjust the operating environment for the library accordingly.

For more information about planning the installation of the TS4500, see this IBM Documentation web page.

2

TS4500 Ultrium Linear Tape-Open and 3592 tape drives

The IBM TS4500 tape library supports Ultrium Linear Tape-Open (LTO) and IBM TS1100 (3592) tape drives.

The high-density 2 (HD2) frames of the TS4500 tape library support the following HD2-compatible tape drive models:

- IBM TS1160 (3592 60E, 3592 60F, and 3592 60S)
- IBM TS1155 (3592 55E and 3592 55F)
- IBM TS1150 (3592 EH8)
- IBM TS1140 (3592 EH7)
- LTO-9 (3588 F9C, 3588 F9S, and 3588 S9C)
- LTO-8 (3588 F8C and 3588 F8S)
- LTO-7 (3588 F7C)
- LTO-6 (3588 F6C)
- LTO-5 (3588 F5C)

This chapter includes the following topics:

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

2.1 IBM TS1100 tape drives for the TS4500 tape library

The TS1100 family of drives for the TS4500 includes the IBM TS1160 tape drive (machine types 3592-60F, 3592-60E, and 3592-60S), IBM TS1155 tape drives (machine types 3592-55F and 3592-55E), IBM TS1150 tape drives (machine type 3592-EH8), and TS1140 tape drives (machine type 3592-EH7). These drives offer a design that is focused on high capacity, performance, and high reliability for storing mission-critical data.

The 3592 family was improved and expanded with the addition of IBM TS1160. The TS1160 is an enhancement over the fifth generation TS1150 and TS1155 tapes drives. It provides unprecedented capacity of 20 TB of uncompressed data on a single tape and new physical host connection options.

> **Note:** This chapter describes only the 3592 models to attach to the TS4500. For more information about other 3592 models, see *IBM Tape Library Guide for Open Systems*, *SG24-5946.*

The 3592 tape drive family features the following common characteristics, which are described in this section:

- Technology enhancements
- Reliability and availability
- Performance or capacity scaling
- Capacity and performance features
- Physical attachment
- Media
- 3592 media cartridge
- WORM functions
- Tape encryption for TS1100

2.1.1 Technology enhancements

The 3592 tape drive family includes the following key features:

- Virtual backhitch, which is the optimum adaptive format and algorithm for improved start-and-stop write performance. For more information, see “Virtual backhitch (nonvolatile caching)” on page 84.
- High performance and robust dual microprocessor architecture. One microprocessor operates the host attachment interface (which is running proven 3590 host attach microcode). The other microprocessor focuses strictly on writing data and reading data from tape. Each microprocessor resets the other microprocessor to act as a fail-safe.
- Statistical Analysis Recording System (SARS) algorithm with extended mount count.
- Fast random access performance when the tape drive operates on any of the Short Length Cartridge (SLC) types.
- Support of an enhanced capacity scaling and segmentation format when the tape drive operates on the full-length, read/write cartridge types JA, JB, JC, JD, and JE, which enable fast locate and read times.
- Streaming lossless data compression (SLDC) algorithm, which is an enhancement of the Lempel-Ziv class 1 (LZ-1) data compression algorithm.
- The JE, JD, JZ, and JL media types contain 16 KB cartridge memory, which is increased from the 8 KB cartridge memory that is contained in JB, JX, JC, JY, and JK media types and 4 KB cartridge memory that is contained in JA, JW, JR, and JJ media types.

Recording format

The IBM 3592 tape drive uses an advanced interleaved bidirectional serpentine recording technique that writes 8, 16, or 32 (depending on the drive) data tracks at a time on a 3592 cartridge. The 3592 cartridge is a half-inch, advanced metal particle, dual-layer tape.

The tape layout consists of five servo bands (prerecorded on the tape) and four data bands where the data is written, as shown in Figure 2-1. The servo bands provide location information to control the positioning of the head as it writes and reads data within the data band. For more information about this design, see “Servo tracks” on page 79.

As shown in Figure 2-1, the area between adjacent servo bands is a *data band*. The 3592 media has four data bands, each with a number of data tracks (128 - 288, which is different for each model).

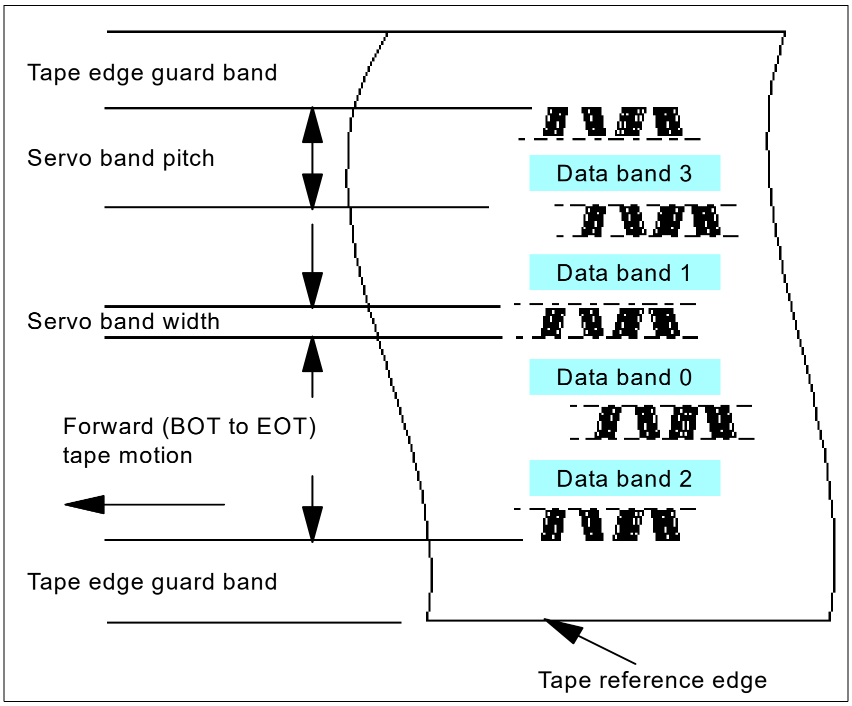

Figure 2-1 Layout of the servo and data bands on the 3592 media

Servo tracks

Servo tracks or bands help to ensure accurate positioning of the tape drive head over the data track so that the head does not stray onto an adjacent track. Servo tracks are necessary to support high-data densities on the tape where the tracks are extremely close together. The servo tracks are written at the time of cartridge manufacture before the cartridge is usable for data storage and retrieval. Each tape write head has two *servo heads*, one servo head for each of the two *servo bands* that it spans.

Two servo bands are used simultaneously to provide two sources of servo information for increased accuracy. Control positions within the servo band are used to reposition the head to write forward and reverse wraps within each of the four data bands. This timing-based servo technology can be finely tuned. This technology supports extremely high-track densities for future 3592 generations because more than eight positions can be defined within the same servo band, which expands the potential track densities.

In addition, significant advances occurred in the tape coating process, which uses high-quality metal particle media.

2.1.2 Reliability and availability

The 3592 tape drive incorporates and expands on the high reliability and function of previous IBM drives that were developed over many years of experience. It builds on proven technologies to enhance and apply new techniques to ensure high reliability and availability.

Improved availability

Improved availability includes the following characteristics:

- Single field-replaceable unit (FRU)

  When a service call is placed, the IBM service support representative (SSR) does not replace any parts or subassemblies inside the canister. The new smaller drive unit means that for any failure within the drive, the IBM SSR exchanges the entire unit rather than performing lengthy diagnostics or component replacement in the field.

- Redundant, hot-pluggable power supplies

  In all configurations, the drives are seated in *cradles* that contain two power supplies. Each pair of power supplies can be used by one or two drives. One power supply is sufficient to run both drives, and the second power supply is provided for redundancy.

- Retention of the Fibre Channel (FC) worldwide name ID during service action

  When a failed drive is exchanged, you do not need to reconfigure the attached hosts or the storage area network (SAN) to recognize a replacement drive. This function also eliminates any issues with SAN hosts finding incorrect addresses during a system restart.

- Retention of the IP configuration for the Ethernet ports during service action

  When a failed drive is exchanged, you do not need to reconfigure the attached Ethernet hosts to recognize a replacement drive. Ethernet port configuration data is maintained during replacement.

When replacing a 3592 drive, the TS4500 maintains all configuration for the new drive by copying the configuration, such as the IP address information and iSCSI names.

The port MAC address might change, depending on which level of LCC card is installed. In the original version of LCC card, the MAC address might have to be changed manually by using the CLI if you want to maintain the same MAC address for each port.

With the latest level of LCC card, the MAC address is maintained the same as the replaced drive; therefore, the MAC address does not need to be set. This card overwrites the MAC address with a TS4500-assigned address. This card is identified by the label on the LCC card (which is shown as Version 4.X), as shown in Figure 2-2.

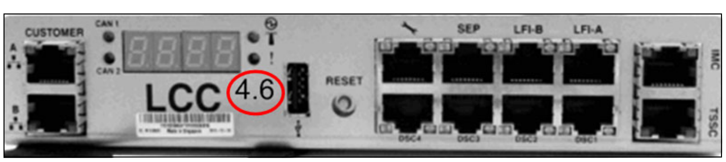

Figure 2-2 New model LCC card

Advanced technology

Advanced technology includes the following characteristics:

- Robust loader mechanism

  The loader mechanism is suitable for the heavy-duty cycle usage in mainframe systems. The leader block on the tape cartridge is replaced by a metal pin, which is enhanced over previous drive implementations for increased robustness.

- Elimination of drive pneumatics and mechanical adjustments

  The aerodynamic movement of the tape over the flat-lap head pulls the tape close to the head while the tape is moving, and provides maximum efficiency in reading and writing. Because of the shape of the head, particles do not accumulate on the tape, which eliminates the possibility of debris contamination of the tape surface. Air-bearing heads effectively cushion the tape that is moving across the head.

  However, whenever the tape stops, it relaxes toward the head surface. The head has a two-stage actuator: one mechanism for moving to the required tape wrap, and another finer actuator for adjustments to the track-following servo.

- Straighter and shorter tape path for better tape tracking

  Tape tracking is improved by using grooved rollers to provide surface-controlled guiding. This enhancement decreases potential wear or damage on the edges of the tape and decreases lateral movement with the shorter tape path.

- Speed matching to reduce backhitching (for more information, see 2.1.3, “Capacity and performance features” on page 82).

  Buffering, speed matching, and virtual backhitch algorithms all serve to eliminate physical backhitching. They improve performance and reduce the wear that is caused by continually braking and reversing direction on the drive mechanics.

- Channel calibration to optimize performance and data integrity

  The drive uses individual read/write data channel calibration, which uses sophisticated techniques that were originally implemented in disk technology.

Enhanced service functions

The following service functions were enhanced:

- Enhanced Statistical Analysis Recording System (SARS) recording

  The tape drive uses SARS to help to isolate the failures between media and hardware. SARS uses the cartridge performance history that is saved in the cartridge memory module and the drive performance history that is kept in the drive flash. The cartridge memory is a serial Electronically Erasable Programmable Read-Only Memory (EEPROM), with read-only and rewritable areas, to determine the more likely cause of failure.

  SARS can cause the drive to request a cleaner cartridge (based on usage) to mark the media as degraded, and to indicate that the hardware is degraded. SARS information is reported through the TapeAlert flags and through media information messages (MIMs) and service information messages (SIMs).

- Diagnostic information

  The drive maintains logs to assist engineering or service personnel. The logs are included in drive memory dumps. They also are accessible to service personnel in several ways. Memory dumps are maintained over Power On Reset (POR).

- More temperature and voltage sensors to improve error isolation

  The drive contains sensors and circuits to detect errors. A temperature sensor monitors the temperature of the drive electronics. Voltage sensors detect when the power supply is out of tolerance. Other error checks, such as tape velocity checks, read/write data integrity checks, and servo checks are performed by using circuitry and sensors. The drive microcode checks for logic errors to handle hardware-detected errors and to detect and report microcode-related errors.

- Drive status indicators and reliability, availability, and serviceability (RAS) functions on the library drive interface

  The drive provides indicators for FC status, whether the power is good, and faults.

- Concurrent microcode update.
- Preventive maintenance

  The 3592 tape drive requires no preventive maintenance beyond the use of the cleaning cartridge. The 3592 media cartridges require proper care and appropriate handling and shipping procedures.

2.1.3 Capacity and performance features

The unique features and specifications of the 3592 make it a true enterprise tape drive in terms of performance and reliability. The following sections describe these industry-leading features.

> **Important:** These features are generic for all 3592 tape drives. The TS1160, TS1155, TS1150, and TS1140 offer other advanced features that are described in 2.2, “IBM TS1160 tape drive” on page 105, 2.3, “IBM TS1155 and TS1150 tape drive” on page 114, and 2.4, “IBM TS1140 tape drive (Model 3592 EH7)” on page 125.

Data buffer

The drive features a large data buffer (1 GB for TS1140, and 2 GB for TS1150, TS1155, and TS1160) with read-ahead buffer management that addresses the lowest band of data rates. It effectively collects more blocks of data in the buffer before it writes to the drive at a higher speed. As a result of this data buffer, the drive stops and starts less often, which generally improves the overall performance and reliability of the drive and tape.

Speed matching

For medium data rates, when the drive is operating from a host that cannot sustain the maximum 3592 data rate, the drive performs dynamic *speed matching*. The drive adjusts the native data rate of the drive as closely as possible to the net host data rate (after it factors out data compressibility). The 3592 drive operates at various speeds (6 - 13 speeds, depending on the drive that is used) when the 3592 format is read or written to in an attempt to match the effective host data rates.

If the net host data rate is between two of the speed matching native data rates, the drive calculates at which of the two data rates to operate. Speed matching reduces the number of required backhitches. In some environments, the backhitch of the drive is masked by the data buffer of the drive. Therefore, the system throughput is not improved or reduced by speed matching.

The following data rate ranges are supported for TS1160 (depending on the logical format and the media type that are used):

- A total of 12 speeds of 122 MBps - 407 MBps for 3592 JE, JV, and JM cartridges that are initialized in J6 format
- A total of 12 speeds of 112 MBps - 365 MBps for 3592 JD, JZ, and JL cartridges that are initialized in J5 and J5A format
- A total of 12 speeds of 99 MBps - 303 MBps for 3592 JC, JY, or JK cartridges that are initialized in J5 format
- A total of 12 speeds of 62 MBps - 252 MBps for 3592 JC or JY cartridges that are initialized in J4 format

The following data rate ranges are supported for TS1155 and TS1150 (depending on the logical format and the media type that are used):

- A total of 12 speeds of 112 MBps - 365 MBps for 3592 JD, JZ, and JL cartridges that are initialized in J5 and J5A format
- A total of 12 speeds of 99 MBps - 303 MBps for 3592 JC, JY, or JK cartridges that are initialized in J5 format
- A total of 12 speeds of 62 MBps - 252 MBps for 3592 JC or JY cartridges that are initialized in J4 format

The following data rate ranges are supported for TS1140 (depending on the logical format and the media type that are used):

- A total of 13 speeds, 76 MBps - 251 MBps for 3592 JC, JK, and JY cartridges that are initialized in Generation 4 format
- A total of 13 speeds, 74 MBps - 203 MBps for 3592 JB or JX cartridges that are initialized in Generation 4 format
- A total of 13 speeds, 41 MBps - 163 MBps for 3592 JB or JX cartridges that are initialized in Generation 3 format
- A total of 13 speeds, 39 MBps - 151 MBps for 3592 JB or JX cartridges that are initialized in Generation 2 format

The 3592 tape drives increase throughput through speed matching while the drive performs the following functions:

- Adjusts tape speed that is based on host data rate
- Calculates the effective host data rate (EHDR)
- Optimizes the data rate by selecting optimal EHDR
- Forces speed changes mid-wrap if advantageous
- Minimizes the time to record data

Cartridge memory

The *cartridge memory* (CM), which is a passive, contactless silicon storage device that is physically a part of the cartridge, is contained within the cartridge. The CM is used to hold information about that specific cartridge, the media in the cartridge, and the data on the media.

The 3592 uses the same CM module as LTO media, with a capacity of 4 KiB for JA, JW, JR, and JJ media, which is extended to 8 KiB on JB, JX, JC, JY, or JK media, and enhanced to 16 KiB for JE, JD, JZ, and JL media. The CM was designed for 3592 to support the high-resolution tape directory feature. The CM differs from the LTO specification because it supports the high-resolution tape directory feature. For more information, see “High-resolution tape directory” on page 83.

Communication between the drive and the CM occurs through a noncontact, passive radio frequency interface (RFI), which eliminates the need for physical connections to the cartridge for power or signals.

High-resolution tape directory

The 3592 drive maintains a tape directory structure with a higher granularity of information about the physical position of data blocks and file marks on the media. This feature gives the 3592 drive improved nominal and average access times for locate operations. Locate times are uniform.

They are based on the position of the block or file mark on the tape, independently of the uniformity of the block size or file mark distribution along the length of the tape. Therefore, the 3592 locate and space performance is targeted to be completely and individually dependent on the longitudinal position on tape of the target block or file mark.

Virtual backhitch (nonvolatile caching)

The 3592 stages write-data through an intermediate dynamic random access memory (DRAM) buffer on its way to tape. This buffer is volatile because it does not retain what is stored if power is lost. For streamed writes (or reads), this buffer yields considerably improved performance.

When the streaming writes cease, a typical pre-3592 tape drive halts the tape and repositions it directly upstream of where the writing ended. From this action, data that is received later can be written immediately after the previously written data. This method eliminates the waste of the considerable length of tape. Substantial lengths of unwritten tape can significantly reduce capacity. Here, a backhitch (reverse) by typical tape drives is used to eliminate this loss of capacity loss after a write to tape.

Nonvolatile caching (NVC) is a 3592 feature that can help greatly improve write performance through backhitch reduction. This system temporarily reserves portions of physical tape for cache areas. Data that is received from the host is written to the volatile buffer as usual, and to nonvolatile tape cache areas, with the exception that no backhitch is necessary when temporary copies are written to cache areas of tape. This temporary capacity loss is easily recouped.

The data is written to temporary cache areas and it is not released in the volatile buffer, but instead it accumulates. This accumulation continues until the buffer is nearly full. Now, the accumulated data in the buffer is rewritten through a streamed write to the standard area of tape.

When the rewrite is complete, the temporary cache areas of tape are released so that they can be overwritten. To significantly improve the average write throughput to tape, temporary copies can be written to the cache areas of tape without backhitching until the buffer is nearly full. Then, a rewrite of the data can be streamed to the standard area of tape.

Aside from the improved write throughput performance, the second effect of NVC writing is to recover the capacity that is lost by the standard writing technique. Data that is received between synchronization events fills containers of data to be written to tape that are called *device blocks* or *data sets*. The standard writing technique calls for padding the last partially filled data set. This padding on average amounts to half the size of the last data set. With the large data set sizes of modern tape drives, this loss can be substantial.

The streaming rewrite of the data that is accumulated in a buffer causes nearly all data sets that are written to a standard area of tape to be written in full, which is known as *data set packing*.

Writing in NVC mode is automatically started by the drive when host writing behaviors are detected that can perform better when in NVC writing mode. Similarly, NVC writing is discontinued when host commands are received that do not benefit from NVC writing, or when commands, such as `Rewind`, are received. When NVC writing is exited, the drive writes any packed data sets that are accumulated in its buffer before it runs the command that stops NVC mode.

Because NVC writing is automatically started and stopped, it is not apparent to host applications. The only indication that NVC writing occurs is the improved capacity and performance that can result from this new mode of writing.

The two components of NVC, backhitch reduction and data set packing, provide major performance and capacity improvements over standard tape drives, such as the 3590, or the Linear Tape-Open (LTO) writing of synchronized data. Data set packing improves overall tape capacity. Backhitch reduction decreases the frequency of mechanical repositions. NVC provides an innovative approach to increasing capacity and write performance in a way that is not apparent to host applications.

Read ahead feature

On sequential reads, the tape drive automatically runs *read ahead* and fills the buffer with data sequentially beyond the target block.

These drives support advanced automatic read-ahead and read-space virtualization at improved access performance and 2x data buffer size. When the drive processes a command to locate or read a block, the drive automatically continues to stream down the tape and read ahead until the data buffer is full. This feature allows subsequent **`Locate`** or **`Read`** commands to be fulfilled from the data buffer at faster speeds, rather than requiring access to the tape.

With this unique function, the drive outperforms competitive drives, which stop and wait for the next command.

2.1.4 Performance or capacity scaling

The 3592 tape drives support scaling and segmentation modes on the read/write (JA, JB, JC, JD, or JE) cartridges so that clients can trade-off capacity for improved access times. Although 256 settings of capacity are supported on the 3592 drive, the following settings are often used:

- Full capacity default mode
- A 20% scaled fast access mode (capacity scaled, front of tape through an x'35' setting)
- Performance scaling for 87% capacity (segmented format, capacity scaling setting x'E0')

Performance scaling, which is also known as *capacity scaling*, is a function through which data can be contained in a specified fraction of the tape, which yields faster locate and read times. This function is possible through the action of modifying internal formatting indicators in the medium and in the cartridge memory chip.

The normal serpentine track format is altered in such a way as to limit the recorded portion of the tape to a specified fraction of the length of the tape, as shown in Figure 2-3. In the 3592, an application can issue a **`Mode Select`** command to scale an individual cartridge. It pertains only to the cartridge that is loaded and it is not persistent.

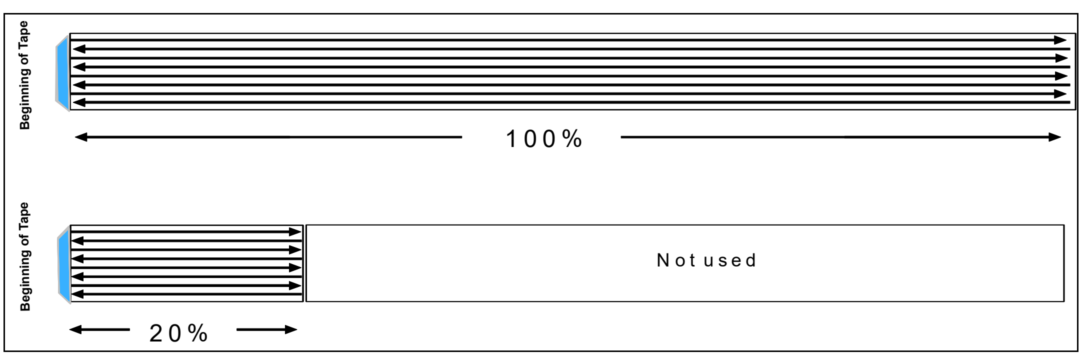

Figure 2-3 Examples for a 100% tape and scaled tape to 20%

The result of performance scaling a tape to a percentage value (for example, 20%) is that the maximum number of recordable gigabytes is reduced to 20% of the normal value. Also, the average time to locate a random record on a full tape that starts from load point is (roughly) 20% of the time to locate a random record from load point for a full, unscaled tape.

The cartridge can be rescaled from any current value to any supported new value. Tape is logically erased by this rescaling. (The end of the data mark is written at the beginning of the tape.) However, tape is not physically erased as with the `long erase` command.

Scaling or rescaling one cartridge does not cause rescaling of the next cartridge. An explicit command must be issued for each cartridge to be scaled or rescaled.

When a scaling operation is requested on a JD type cartridge that uses the TS1155 drive, the media is up-formatted to the J5A logical format at the same time the scaling operation is performed unless the format is controlled through explicit means.

Performance segmentation

Performance segmentation provides fast access and capacity by allowing the tape to be divided into two segments. One segment is a fast access segment to be filled first, and the other segment is more capacity to be filled after the first segment. Therefore, it is high performance in two ways. It has segmentation and high-performance random access in the first segment, as though it was a scaled cartridge, while it provides other larger capacity, as shown in Figure 2-4.

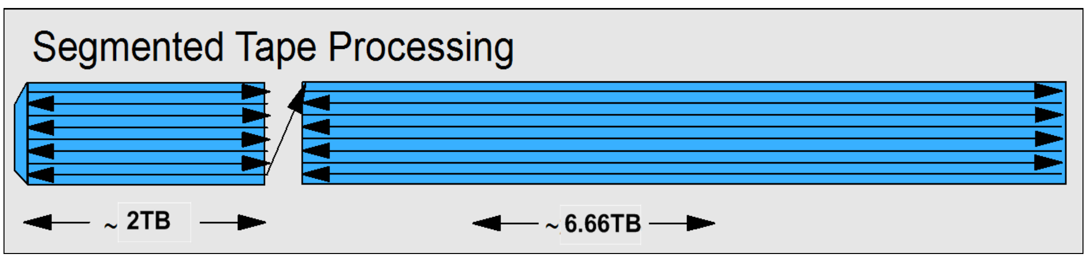

Figure 2-4 Segmented tape on 3592 E08 JD media

Performance capacity scaling and segmentation have the following implications:

- If host systems provide a means to limit the amount of data that a client places on the media, for example, with a percent usage construct, the user achieves a much faster average access time to the first data. Also, more locates on the same volume improve significantly.
- With segmentation, a less than 1% degradation occurs in the data rate because of the increased number of wrap changes. Segmentation also reduces the nominal cartridge capacity by approximately 10%.

Important: Capacity scaling is not supported for economy or WORM tapes.

For more information about for the capacity scaling limitations and instructions for setting up, see *IBM Tape Device Drivers Installation and User's Guide,* GC27-2130.

Access performance specifications and drive characteristics

The access performance and drive characteristics of the 3592-60E, 3592-60F, 3592-60S, 3592-55E, 3592-55F, 3592-EH8, and 3592-EH7 are listed in Table 2-1 on page 87. For better performance, the block size of the data should be 256 K or greater.

Table 2-1 Access performance specifications and drive characteristics

Parameter	3592-60E, 60F, and 60S	3592-EH8, 55E, and 55F	3592-EH7
Tape speed, locate/rewind	12.4 mps	12.4 mps	12.4 mps
Drive load/ready time	12 seconds (s)	12 seconds (s)	15 s
Block locate time from load point average	45 s for JE and JV 13 s 20% scaled JE 40 s for JC and JY 45 s for JD and JZ 11 s for JK 13 s for JL 12 s 20% scaled JC 13 s 20% scaled JD	40 s for JC and JY 45 s for JD and JZ 11 s for JK 13 s for JL 12 s 20% scaled JC 13 s 20% scaled JD	37 s for JB and JX 40 s for JC and JY 11 s for JK 15 s 20% scaled JB 12 s 20% scaled JC
Time to first data average (load/ready + locate)	55 s for JE and JV 23 s 20% scaled JE 50 s for JC and JY 55 s for JD and JZ 22 s for JK 23 s for JL 23 s 20% scaled JC 23 s 20% scaled JD	50 s for JC and JY 55 s for JD and JZ 22 s for JK 23 s for JL 23 s 20% scaled JC 23 s 20% scaled JD	42 s for JB and JX 55 s for JC and JY 26 s for JK 30 s for 20% scaled JB 27 s for 20% scaled JC
Unload time	31 s for JE, JV, JM, JD, JZ, and JL 36 s for JC, JY, and JK	31 s for JD, JZ, and JL 36 s for JC, JY, and JK	24 s for JB and JX 36 s for JC, JY, and JK
Maximum rewind time	94 s 100% scales JE, and JV 34 s 20% scaled JE, and JV 76 s 100% scaled JC and JY 26 s 20% scaled JC 18 s JK 94 s 100% scaled JD and JZ 34 s 20% scaled JD and JZ 34 s JL-	76 s 100% scaled JC and JY 26 s 20% scaled JC 18 s JK 94 s 100% scaled JD and JZ 34 s 20% scaled JD and JZ 34 s JL	72 s 100% scaled JB and JX 24 s 20% scaled JB 76 s 100% scaled JC and JY 26 s 20% scaled JC 18 s JK
Native data rate	400 MBps	360 MBps	250 MBps
Device data rate: Maximum that is sustained with maximally compressible data	900 MBps	700 MBps 600 Mbps for 3592-55E	650 MBps
Interface burst transfer rate: Maximum	1600 MBps (FC-16)	800 MBps (FC-8)	800 MBps (FC-8)
Number of tracks	J6 format, 8704 JE, JV, and JM J6A format, 7680 JD J6 format, 5120 JD, JZ, and JL J6 format, 4608 JC, JK, and JY	J5A format, 7680 JD J5 format, 5120 JD, JZ, and JL J5 format, 4608 JC, JK, and JY	J4 format, 1792 JB, and JX J4 format, 2560 JC, JK, and JY
Number of passes (from BOT to EOT)	J6 format, 272 JE, JV, and JM J6A format 240 JD J6 format 160 JD, JZ, and JL J6 format 144 JC, JK, and JY	J5A format 249 JD J5 format 160 JD, JZ, and JL J5 format 144 JC, JK, and JY	J4 format 56 JB and JX J4 format 80 JC, JK, and JY
Linear density	555 kilobits per inch (Kbpi)	510 kilobits per inch (Kbpi)	500 Kbpi
Servo regions	5	5	5
Data tracks recorded simultaneously	32	32	32
Buffer size	2 GB	2 GB	1 GB

2.1.5 Physical attachment

The 3592 (model EH7, EH8, 55F, 60F, and 60S) tape drives are supported on several tested operating systems and platforms. For more information about which systems were tested and are approved for use, see this web page of the IBM System Storage Interoperation Center (SSIC).

The 3592 (model 55E and 60E model) tape drives are supported on open systems, servers that are running Microsoft Windows only. Approval of i-RPQ 8B3685 is required to order a TS1155 Tape Drive Model 55E and TS1160 Tape Drive Model 60E.

The 3592 (model EH7, EH8, 55E, 55F, 60E, 60F, and 60S) are supported for attachment in the IBM TS4500, and communicates with the TS4500 tape library through an internal Ethernet interface. It uses Statistical Analysis and Reporting System (SARS) to isolate failures between the media and hardware.

The 3592 model EH7, EH8, 55F, and 60F offer a dual-port 8 Gbps or 16 Gbps Fibre Channel host attachment interface. This feature provides flexibility in open systems environments because drives can attach to open systems servers directly with Fibre Channel attachments.

The 3592 models 55E and 60E tape drive have a dual-ported 10 or 25 Gb optical Ethernet ports for host attachment. This drive was optimized for cloud-based and large, open-compute environments.

The 3592 model 60S tape drive provides a dual-port 12 Gb SAS (Serial Attached SCSI) interface for host attachment. This drive brings more versatility to businesses with substantial storage, backup, and archiving demands with a cost-competitive communications interface to help simplify storage management and system performance.

These drives feature similar back and front panels, which are shown in Figure 2-5 and Figure 2-6 on page 90.

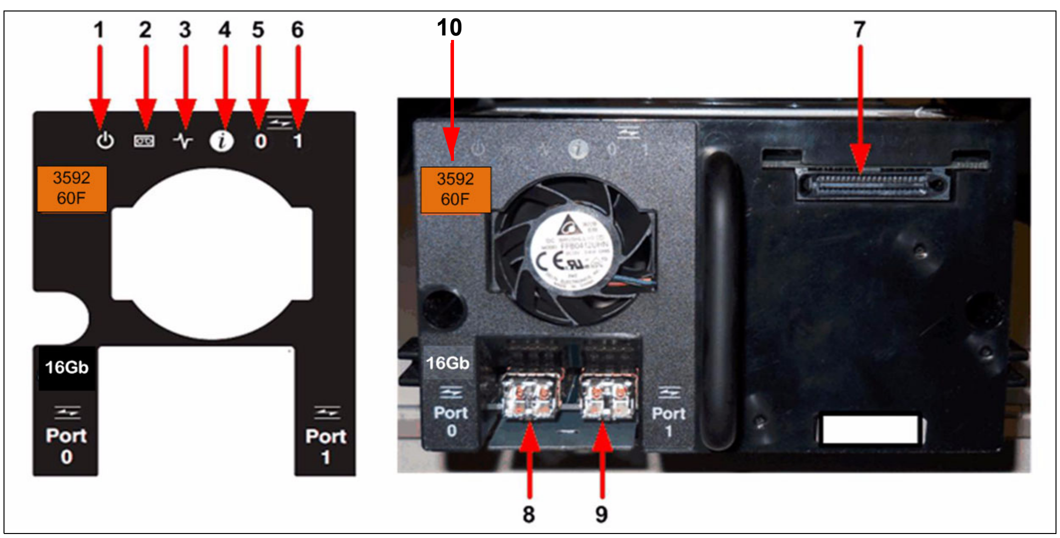

Figure 2-5 TS4500 3592 drive rear panel

All of the rear panel components and the normal status of the light-emitting diodes (LEDs) are listed in Table 2-2 on page 89. The numbers in Table 2-2 on page 89 correlate to the numbers that are shown in Figure 2-5 rear panel.

Table 2-2 TS4500 rear panel components

Number	Description	LED color or description
1	Power status	Green
2	Drive status	Yellow or green
3	Library connection	Green
4	Information	Blue, possible canister card issue. Drive fan starts running at high speeds and turns on blue LED.
5	Host port 0 activity	Green or yellow
6	Host port 1 activity	Green or yellow
7	TS4500 interface connector	TS4500 communication and power
8	Port 0, 8, or 16 Gbps Fibre Connection Models EH7, EH8, 55F, and 60F	Supports N_ports and NL ports only, direct Fabric attach protocol
9	Port 1, 8, or 16 Gbps Fibre Connection Models EH7, EH8, 55F, and 60F	Supports N_ports and NL_ports only, direct Fabric attach protocol
8	Port 0, 10, or 25 Gbps Ethernet connection Models 55E and 60E	RoCE v2 protocol. Short wave multimode optical SFP transceiver
9	Port 1, 10, or 25 Gbps Ethernet connection Models 55E and 60E	RoCE v2 protocol. Short wave multimode optical SFP transceiver
8	Port 0, 12 Gbps SAS connection Model 60S	SAS (Serial Attached SCSI) attachment interface
9	Port 1, 12 Gbps SAS connection Model 60S	SAS (Serial Attached SCSI) attachment interface
10	Drive Model	Drive model 3592 EH7, EH8, 55E, 55F, 60E, 60F, or 60S

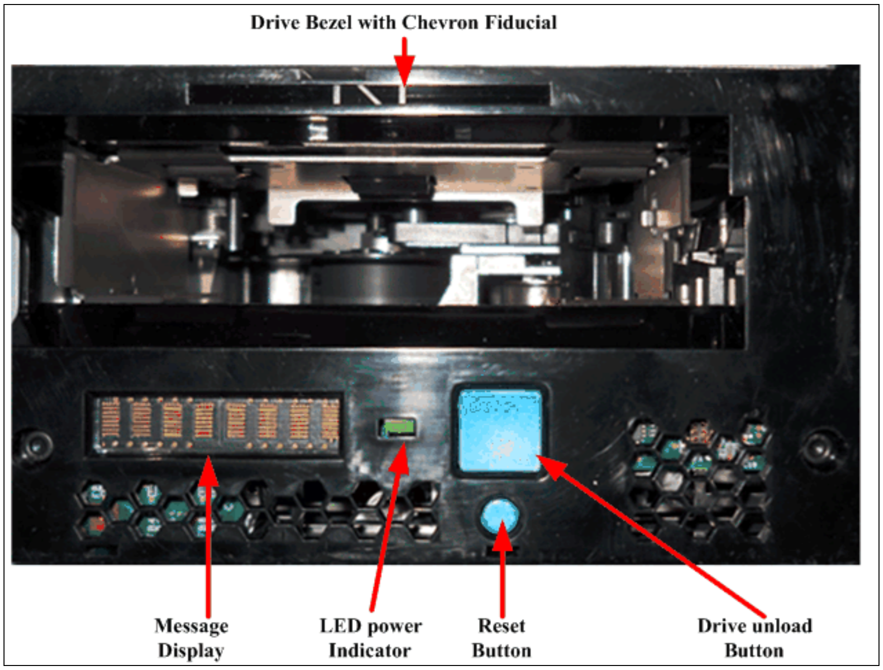

Figure 2-6 TS4500 3592 drive front panel

Fibre Channel drives

The TS4500 3592 Fibre Channel (FC) tape drives includes dual-ported 8 or 16 Gbps, switched FC attachments, which provides attachment to multiple servers or a single server with redundancy. This function also offers attachment flexibility in an open systems environment. The drive can be directly attached to open systems servers with use of a Fibre Channel host bus adapters (HBAs).

The TS1160 3592 60F tape drive attempt to connect at 16 Gbps. However, they autonegotiate down to 8 Gbps, or 4 Gbps if the system or port that they are connected to cannot support higher bandwidth. The 3592 55F, EH8, and EH7 tape drives attempt to connect at 8 Gbps. However, they can autonegotiate down to 4 Gbps, 2 Gbps, or 1 Gbps if the system or port they are connected to cannot support higher bandwidth.

The 3592 8 Gb Fibre Channel (FC) attached tape drives can operate as a node loop port (NL_port) (Fibre Channel Arbitrated Loop [FC-AL] support) or as a node (N_port) (supporting direct connection to a SAN switch, which is also known as *point-to-point* or fabric mode). The 3592 tape drives autonegotiate to the N_port or NL_port, depending on whether a loop or a point-to-point connection is detected when the drive boots.

The drives do not autonegotiate, if the drive was set to use an explicit setting of these configurations. Regardless of whether the 3592 tape drives connect as an NL_port or an N_port, they autonegotiate to be a public device (attached to a switch) or a private device (attached to another N_port; that is, directly attached to a host).

The 3592 16 Gb Fibre Channel (FC) connected tape drives operate in node or fabric mode (N_port) only, supporting direct connection to a SAN switch or supported HBA. This is also known as *point-to-point* or fabric mode. This port does not support FC-AL (Fibre Channel Arbitrated Loop) mode.

If a library drive is replaced, an IBM SSR selects the replacement unit to automatically inherit the configuration attributes of the failed unit. This way, a user can avoid reconfiguring the zoning in the switches. Alternatively, the Management Interface (MI) can be used to change these fields directly at any time.

For more information about FC attachment planning, see *IBM TotalStorage Enterprise Tape System 3592 Introduction and Planning Guide,* GA32-0555.

For more information about applications and their levels that support 3592 tape drives, see the ISV matrix, which is available to download as a PDF file from the *Independent Software Vendor Matrix (ISV) for IBM TotalStorage 3592 tape drives and LTO*.

Multiple Fibre Channel ports

All FC 3592 models have two independent FC interfaces or ports. Both ports run the Small Computer System Interface (SCSI) protocol with FC tape support. By using these two ports, concurrent attachment of two independent FC configurations can be made to each drive. One or both ports can be attached to various open systems servers, SAN switches, and directors.

The 3592 tape drives support industry-standard shortwave LC-Duplex fiber optic cables, with cable lengths of up to 500 m (1,640 ft.) and 50 microns of core fiber, depending on the required attachment speed.

The following maximum distances are supported by shortwave adapters with the 50/125 LC-Duplex fiber optic cables:

- 1 Gbps: 500 meters (1,640 feet)
- 2 Gbps: 300 meters (984 feet)
- 4 Gbps: 150 meters (492 feet)
- 8 Gbps: 50 meters (164 feet)
- 16 Gbps: 35 meters (115 feet)

The following maximum distances are supported by shortwave adapters with the OM3 LC-Duplex fiber optic cables:

- 8 Gbps shortwave adapters have a maximum distance of 150 meters (492 ft.)
- 16 Gbps shortwave adapters have a maximum distance of 100 meters (328 ft.)

Supported topologies

The 3592 tape drives support switched fabric and point-to-point loop topologies.

Switched fabric

Two or more FC endpoints connect through a switch. The FC architecture supports up to 256 ports through each switch. Switches include a function that is called *zoning*. By using this function, you can partition the switch ports into port groups and then assign group access to other groups. This function prevents group interference. With switched fabrics, all of their ports have simultaneous use of the full FC architecture bandwidth.

Point-to-point loop

A point-to-point loop is similar to a point-to-point topology. Both have two connected FC endpoints. The difference is in the protocol. Therefore, when only two FC endpoints are connected, either protocol is usable. However, both endpoints must use the same protocol. The 3592 model supports a point-to-point loop. Most FC adapters default to the loop protocol when they are not directly connected to a fabric.

Important: SAN switches normally default the switch port to loop mode. If the port is set to automatic mode, loop mode is the first mode that is attempted during the port login process. The 3592 accepts loop mode and logs in to the port. To get the 3592 to log in to the SAN switch port in fabric mode, the port in the switch must be set to fixed fabric mode by the switch administrator or in the drive by using the management interface. The TS1160 model 60F 16 Gb FC ports do not support FC-AL.

Address assignments

The 3592 tape drives must have an FC address to communicate over the FC interface. The tape drives support hard and soft addressing. Most FC hosts (initiators) support hard addressing and do not support soft addressing. For more information, see the device driver documentation.

Fibre Channel worldwide name ID

Each Fibre Channel card on the 3592 tape drive has four names (Node 0, Node 1, Port 0, and Port 1) that are hardcoded into the electronics of the card by IBM manufacturing. These names are similar to a serial number and are unique any place in the world. The TS4500 configures a World Wide Name (WWN) for each drive. The WWN is derived from the location of the drive that is inside the TS4500 and from the serial number and physical location of the TS4500.

More information about SAN best practices and IBM SAN switches, see *IBM Tape Library Guide for Open Systems*, SG24-5946.

Ethernet drives

IBM TS1155 Tape Drive, Model 55E, delivers 10 Gb, and the IBM TS1160 Tape Drive, Model 60E, delivers 10 Gb or 25 Gb Ethernet host attachment interface, which is optimized for cloud-based and hyperscale environments.

This configuration provides dual 10 Gb or dual 25 Gb optical Ethernet host attachment ports that use Remote Direct Memory Access (RDMA) over Converged Ethernet for cloud-based and open-compute environments.

The dual 10 Gb and dual 25 Gb ports uses shortwave multi-mode optical SFP transceivers and the allowable cable lengths are listed in Table 2-3.

Table 2-3 Ethernet optical cable limits

Fiber cable type	Connector Type	Minimum modal bandwidth at 850 nm (MHz x km)	Operating range (in meters)
10 Gb 62.5 μm MMF	LC	160	2 - 26
10 Gb 62.5 μm MMF	LC	200	2 - 33
10 Gb 50 μm MMF	LC	400	2 - 66
10 Gb 50 μm MMF	LC	500	2 - 82
10 Gb 50 μm MMF	LC	2000	2 - 300
16 Gb 50 μm MMF	LC	2000	2 - 300
16 Gb 50 μm MMF	LC	4700	2 - 400

The TS1100 Ethernet drives currently supports iSCSI Extension for RDMA (iSER) on Converged Ethernet (RoCEv2). This specific protocol uses a UDP transport layer and required Data Center Bridging (DCB) switches and lossless networks.

Extensions for RDMA (iSER) is a standard that enables iSCSI hosts and targets to take advantage of RDMA capabilities. iSER runs on top of an RDMA capable Network Interface Card (rNIC) regardless of the protocol.

The TS1155 55E and TS1160 60E are supported by the Microsoft Windows device driver and requires approval of i-RPQ 8B3685.

Serial Attached SCSI (SAS) drives

IBM TS1160 Tape Drive, Model 60S includes a dual-port 12 Gb SAS interface that brings more versatility to businesses with substantial storage, backup, and archiving demands with a cost-competitive communications interface to help simplify storage management and system performance.

A Mini-SAS HD to Mini-SAS HD cable is required for attaching a TS1160 model 60S tape drive to host adapters or other storage area network components.

TS1160 Model 60S offers 12 Gb speed of connectivity; however, they can autonegotiate down to 6 Gb or 3 Gb if the system or port they are connected to cannot support higher bandwidth.

Note: Copper SAS cables longer than 3 meters are not supported.

2.1.6 Media

Users must cost-effectively store more digital information than ever before, often to meet growing regulatory and legal requirements. The 3592 tape drives help meet these needs with IBM Tape Cartridge 3592. The TS1160, TS1155, TS1150, and TS1140 all use the 3592 tape cartridge. This tape cartridge offers various capacity options, depending on the drive and the recording format that is used or the cartridge model that was ordered: Data, WORM, or Economy.

These capabilities expand the range of client data workloads that can be addressed with the 3592 tape drives. The economy cartridge can help lower the cartridge cost for users with smaller capacity needs and provide faster access to data. The WORM cartridges provide nonerasable, nonrewritable storage media. Users with regulatory or legal requirements to store electronic records for long periods might be able to use the 3592 tape drives to provide cost-effective storage.

The 3592 cartridges have a form factor that is similar to the 3590 tape cartridge. They are supported in the following IBM cartridge library environments:

- IBM TS3500
- IBM TS4500

The IBM 3592 half-inch tape cartridge contains an advanced fourth-generation metal particle formulation in a dual-layer coating on a half-inch-wide tape. The IBM tape uses an advanced magnetic coating and process that provides a high output and signal quality to support the current 3592 tape drives.

The tape features an ultra-smooth and uniform magnetic layer that is less than 0.2 microns thick and a specially refined coating formulation that is designed to help improve media reliability and performance and minimize the wear of the tape heads and components. A precision timing-based servo with enhanced features helps enable high track densities, high data rates, data access performance, high reliability, and stop-start performance.

The following media are used for the different media types:

- Dual-coat, MP nanocubic particle, and Polyethylene naphthalate (PEN) substrate 8.9 µm nominal thickness (JA types)
- Dual-coat, MP nanocubic particle, and PEN substrate 6.6 µm nominal thickness (JB types)
- Dual-coat, Barium Ferrite (BaFe) particle, PEN substrate, and 6.1 µm nominal thickness (JC types)
- Dual-coat, BaFe particle, Aramid substrate, and 5.0 µm nominal thickness (JD types)
- Hc perpendicularly oriented BaFe particle (JE types)

Modifications to the cartridge design and construction help improve pin retention, hub and clutch engagement, spool alignment, and tape stacking within the cartridge. These enhancements help improve reliability and durability of the media and the tape drive. Enhanced assembly strengthens the cartridge at critical locations and helps make the 3592 cartridge less susceptible to damage, such as damage from being dropped.

The tape is pulled from the cartridge with a leader pin rather than a leader block as in the 3590 cartridge. A sliding door covers the area that was formerly occupied by the leader block in a 3590 cartridge. A locking mechanism prevents the media from unwinding when the cartridge is not in a drive. A special mechanical design provision prevents the 3592 cartridge types from being loaded into 3590 or 3490 drives. If a 3592 cartridge is inadvertently loaded into a 3590, the cartridge present sensor does not change state and the drive does not attempt to load.

2.1.7 3592 media cartridge

This section provides more information about the 3592 cartridge media that are supported by the TS4500 tape library.

Media types and compatibility

All 3592 tape drives support cartridge reuse. The 3592 tape cartridges can be reformatted to any tape format supported by the tape drive when it writes from BOT. When reformatting, all data on the cartridge is erased.

> **Note:** Cartridge reuse depends on the compatibility of the media on the drive that is used.

The cartridges that are compatible with the 3592 tape drives that are supported by the tape library are listed in Table 2-4.

Table 2-4 Cartridges that are compatible with 3592 tape drives

Media	Recording format	TS1160 tape drive		TS1155 tape drive		TS1150 tape drive		TS1140 tape drive	
		Encryption enabled	Encryption not enabled	Encryption enabled	Encryption not enabled	Encryption enabled	Encryption not enabled	Encryption enabled	Encryption not enabled
JE, JM, JV[a]	Encrypted 60F	RW	Reformat	-	-	-	-	-	-
	60F	RW	RW	-	-	-	-	-	-
JD, JL, JZ[a]	Encrypted 55F	RW	Reformat	RW	Reformat	-	-	-	-
	55F	RW	RW	RW	RW	-	-	-	-
	Encrypted E08	RW	Reformat	RW	Reformat	RW	Reformat	-	-
	E08	RW	RW	RW	RW	RW	RW	-	-
JC, JK, JY[a]	Encrypted E08	RW	Reformat	RW	Reformat	RW	Reformat	-	-
	E08	RW	RW	RW	RW	RW	RW	-	
	Encrypted E07	RO	Reformat	RO	Reformat	RW	Reformat	RW	Reformat
	E07	RO	RO	RO	RO	RW	RW	RW	RW
JB, JX[a]	Encrypted E07	-	-	-	-	-	-	RW	Reformat
	E07	-	-	-	-	-	-	RW	RW
	Encrypted E06	-	-	-	-	-	-	RW	Reformat
	E06	-	-	-	-	-	-	RW	RW
	Encrypted E05	-	-	-	-	-	-	RO	Reformat
	E05	-	-	-	-	-	-	RO	RO
JA	Encrypted E06	-	-	-	-	-	-	RO	Reformat
	E06	-	-	-	-	-	-	RO	RO
	Encrypted E05	-	-	-	-	-	-	RO	Reformat
	E05	-	-	-	-	-	-	RO	RO
	J1A	-	-	-	-	-	-	RO	RO

a. WORM tapes can be reformatted only if they are initialized and labeled. If data exists on the tape, reformatting is not allowed. Labels: (-) Not supported, (RW) Read Write, and (RO) Read Only

Two basic formats are used: Enterprise Format (EFMT) and Enterprise Encrypted Format (EEFMT). Each tape drive model has different formatting capabilities for the TS4500 compatible drives, as listed in Table 2-5.

Table 2-5 Read and write format types that are supported

Drive type	EFMT1	EFMT2 EEFMT2	EFMT3 EEFMT3	EFMT4 EEFMT4	EFMT5 EEFMT5	EFMT6 EEFMT6
TS1160	No	No	No	Read only	Read/write	Read/write
TS1155 TS1150	No	No	No	Read/write	Read/write	No
TS1140	Read only	Read only	Read/write	Read/write	No	No

Note: The TS1160 and TS1155 are not supported for z/OS attachment.

By using their supported densities, 3592 tape drives can use different media. The capability of each drive to use media inside a TS4500 is listed in Table 2-6.

Table 2-6 Drive and cartridge compatibility

Tape unit	JE, JM, JV cartridge	JD, JL, or JZ cartridge	JC, JK, or JY cartridge	JB or JX cartridge
TS1160	Read/write	Read/write	Read/write	No
TS1150, TS1155	No	Read/write	Read/write	No
TS1140	No	No	Read/write	Read/write

The media types, native capacity options, and compatibility options that are available with 3592 tape drives that are supported by the TS4500 tape library are listed in Table 2-7.

Table 2-7 IBM Enterprise 3592 media types

Media description	Media type	3592 60F, 60E, 60S format J6 native capacity	3592 55E, 55F format J5A native capacity	3592 EH8 format J5 native capacity	3592 EH7 format J4 native capacity
Extended data	JB	Not supported	Not supported	Not supported	1.6 TB (1.46 TiB)
					1 TB (.9 TiB) E06 format
Advanced data	JC	7 TB (6.37 TiB)	7 TB (6.37 TiB)	7 TB (6.37 TiB)	4 TB (3.6 TiB)
Advanced data	JD	15 TB (13.64 TiB)	15 TB (13.64 TiB)	10 TB (9.1 TiB)	Not supported
Advanced data	JE	20 TB (18.19 TiB)	Not supported	Not supported	Not supported
Advanced economy	JK	900 GB (.82 TiB)	900 GB (.82 TiB)	900 GB (.82 TiB)	500 GB (.45 TiB)
Advanced economy	JL	3 TB (2.73 TiB)	3 TB (2.73 TiB)	2 TB (1.82 TiB)	Not supported
Advanced economy	JM	5 TB (4.55 TiB)	Not supported	Not supported	Not supported
Extended WORM	JX	Not supported	Not supported	Not supported	1.6 TB (1.46 TiB)
					1 TB (.9 TiB) E06 format
Advanced WORM	JV	20 TB (18.19 TiB)	Not supported	Not supported	Not supported
Advanced WORM	JY	7 TB (9.1 TiB)	7 TB (6.37 TiB)	7 TB (6.37 TiB)	4 TB (3.6 TiB)
Advanced WORM	JZ	15 TB (13.64 TiB)	15 TB (13.64 TiB)	10 TB (9.1 TiB)	Not supported

Important: The TS1160, TS1155, and TS1150 cannot read or write to JA, JW, JJ, JR, JB, or JX media.

Media types JA, JW, JR, and JJ are supported as read-only by media types JB and JX with E05 format are supported as read-only by the 3592-EH7 drive.

Figure 2-7 shows an example of the media types: a full length read/write tape on the left, WORM cartridges in the middle, and economy read/write cartridges on the right. The WORM cartridges have a platinum-colored shell, and the read/write cartridges have a black shell.

Figure 2-7 IBM Enterprise 3592 WORM and read/write cartridges

Labels

The 3592 cartridges use a media label to describe the cartridge type. Figure 2-8 shows a 3592 JE cartridge label. In tape libraries, the library vision system identifies the types of cartridges during an inventory operation. The vision system reads a *volume serial number* (VOLSER), which is on the label on the edge of the cartridge. The VOLSER contains 1 - 6 characters, which are left-aligned on the label. If fewer than 6 characters are used, spaces are added. The media type is indicated by the seventh and eighth characters.

J A G 1 2 3 J E

Figure 2-8 View of the 3592 cartridge label

Cleaning cartridges

One cleaning cartridge is designed specifically for the 3592 tape drives. As with the data cartridges, the 3592 cleaning cartridges are not interchangeable with any other model cleaning cartridges (for example, LTO cleaning cartridges). Therefore, both types of cleaning cartridges must be inserted into the library if both types of drives are in the environment.

The cleaning cartridge also contains a cartridge memory device, which automatically tracks the number of times that it was used. Cleaning cartridges must be replaced after 50 cleaning cycles.

The physical characteristics of the 3592 cleaning cartridge can be used to distinguish it from the 3592 data cartridges. The product label on the top of the cartridge is white with the word "cleaning" printed on it. Instead of the write-protect switch, a non-movable light gray block exists, which is shown as number 1 in Figure 2-9 on page 99.

The cartridge door is also light gray. If you order cleaning cartridges with pre-attached labels, the first three characters of the VOLSER are CLN, as identified by number 2 in Figure 2-9.

Figure 2-9 Cleaning cartridge

> **Note:** The IBM Enterprise Tape Cartridge is universal. It can be used on any model 3592 tape drive.

2.1.8 WORM functions

All 3592 tape drives with the appropriate installed microcode version can read and write WORM cartridges. The WORM data cartridges for the IBM 3592 tape drive provide nonalterable, nonrewritable tape media for long-term records retention. The WORM cartridges include the following characteristics:

- WORM cartridges are available in the following formats for the TS4500:
  - JV (advanced), which is supported by the 3592 60E, 60F, and 60S.
  - JZ (advanced), which is supported by the 3592 EH8, 55E, 55F, 60E, 60F, and 60S with 55F format, and by the 55E, 55F, 60E, 60F, and 60S with 60F format.
  - JY (advanced), which is supported by the 3592 EH7 tape drive with EH7 format, and by the 3592 EH8, 55E, 55F, 60E, 60F, and 60S with EH8 format.
  - JX (extended), which is supported by the 3592 EH7 in EH7 and EH6 formats.
- Non-reversible screws are used to secure the media housing.
- WORM and read/write cartridges can be intermixed in the same IBM TS4500 tape library.
- When the drive senses that a cartridge is a WORM cartridge, the microcode prohibits changing or altering user data that is already written on the tape. The microcode tracks the last point on the tape to which data can be appended with an overwrite-protection pointer that is stored in the cartridge memory.
- Each WORM cartridge is identified by using a Worldwide Cartridge Identifier (WWCID), which is permanent and locked. The WWCID provides another level of security for data that must be maintained.
- A WORM cartridge can never be changed to non-WORM, and a non-WORM cartridge cannot be changed to WORM.
- User data that is written on WORM cannot be modified, reformatted, or erased.

WORM basics

The 3592 tape drives support 3592 read/write cartridges and 3592 WORM cartridges. The WORM cartridge is geometrically identical to a read/write cartridge. It uses the same rewritable media formulation. However, the servo format, which is mastered onto the tape at manufacturing, is different for WORM cartridge types.

The WORM function does not come from any inherent non-reversible media characteristic (such as permanent WORM on optical CD-R media or optical WORM). Instead, the WORM function is enabled by the method by which the microcode of the 3592 drive handles a WORM cartridge.

The microcode of the drive does not support the overwrite or erasure of previously written user data, such as records or file marks. However, the microcode of the drive supports appending new data after the existing data.

Unique cartridge identifier

Each IBM 3592 tape WORM cartridge is identifiable through a unique cartridge identifier (UCID). The intent of the UCID is that it is constructed to ensure that it is unique worldwide. This identifier is derived from the 4-byte unique cartridge memory serial number of the cartridge memory chip in the 3592 WORM cartridge.

This serial number is concatenated with the 8-byte unique tape serial number that was created from information that was mastered into the timing-based servo at the time that the cartridge was manufactured.

The parts of UCID that come from this combined serial number are written to a locked part of the cartridge memory. This other level of security supports legal audit requirements. Furthermore, the UCID supports unique cartridge tracking, and it can be the differentiator to using other WORM tape providers.

Drive operation to prevent overwriting

A WORM drive handles a WORM cartridge differently than a read/write cartridge. In general, a WORM drive responds to a subset of the Small Computer System Interface (SCSI) commands that work on a read/write cartridge. For example, an **`Erase`** command is rejected with the correct error posted.

Additionally, a WORM drive rejects certain command sequences of otherwise valid commands. For example, if a cartridge is not empty, a **`Rewind`** command that is followed by a **`Write`** command is rejected with the correct error posted.

The microcode tracks the last point that can be appended on the tape by using an overwrite-protection pointer that is stored in the CM. The SARS data can be written and updated on WORM tapes because the SARS data is not in the user area of the tape.

The 3592 tape drives allow append operations to data that is already on WORM cartridges, and allow overwriting of file marks and other non-data attributes to provide application transparency. However, they do not allow overwriting data under any circumstances. After they are full of data, WORM cartridges cannot be reused or erased by the drive, and they must be physically destroyed or bulk-degaussed to delete data. For full tape application use, certain trailer and label record overwrites are allowed.

> **Important:** Because WORM cartridges cannot be reused after they are written to, WORM cartridges must be physically destroyed when they are no longer of use. If the WORM cartridge includes sensitive data, it must be bulk-erased before it is discarded. This process erases everything on the tape, including the mastered servo pattern, which renders the tape useless.

2.1.9 Improved media Statistical Analysis and Reporting System

The 3592 drives support Statistical Analysis and Reporting System (SARS).

The tape drive uses SARS to help isolate failures between media and hardware. SARS uses the cartridge performance history, which is saved in the cartridge memory (CM) module, and the drive performance history, which is kept in the drive flash electrically erasable programmable read only memory (EEPROM), to determine the likely cause of failure. SARS can cause the drive to request a cleaning tape to mark the media as degraded, and to indicate that the hardware is degraded.

SARS information is reported through the TapeAlert flags and through media information messages (MIMs) or service information messages (SIMs).

The 3592 drive maintains a history of the last 100 mounts for both Volume Statistical Analysis and Reporting System (VSARS) and Hardware Statistical Analysis and Reporting System (HSARS).

> **Note:** Media SARS information is preserved when media is reformatted.

Starting the TS1150 implements an enhanced SARS function that is known as client-centric SARS (ccSARS). This function improves the overall amount of information that is maintained, and the presentation means to the client in concert with the automation system.

The media SARS function for the drives includes the following actions:

- Tape alerts are generated when media passes usage life, as determined by full-file passes, meters of tape that were processed, or the write pass count, and the total number of mounts (which was already supported).
- A media SARS summary is maintained in the cartridge memory in a manner where it can be rebuilt on tape if the SARS records on tape cannot be read and must be reinitialized. This cartridge memory copy is also readable on an earlier level TS1140 drive to preserve SARS information between logical format conversions.

2.1.10 Tape encryption for TS1100

The TS1160, TS1155, TS1150, and TS1140 tape drives use an Advanced Encryption Standard (AES) encryption key, which is a random string of bits that is generated specifically to scramble and unscramble data. Encryption keys are created by using algorithms that ensure that each key is unique and unpredictable. The longer the key string is, the harder it is to break the encryption code. These drives use 256-bit AES algorithm keys to encrypt data.

The following types of encryption algorithms are used for encryption:

- Symmetric algorithms

  Symmetric (or *secret key*) encryption uses a single key for encryption and decryption. Symmetric key encryption generally is used for encrypting large amounts of data efficiently.

- Asymmetric algorithms

  Asymmetric encryption uses a pair of keys. Data that is encrypted by using one key can be decrypted only by using the other key in the asymmetric key pair.

When an asymmetric or public/private key pair is generated, the public key is used for encryption, and the private key is used for decryption.

TS1100 family of tape drives in the TS4500 use both types of encryption algorithm. Symmetric encryption is used for high-speed encryption of user or host data. Asymmetric encryption (which is slower) is used to protect the symmetric key that is used to encrypt the data (key wrapping).

The TS1160, TS1155, and TS1150 support the capability to record encrypted and plain data on a volume, under the following circumstances:

- The drive must be in Application Managed Encryption (AME) (T10 method).
- In the T10 mode, encryption is controlled on a block-by-block basis by the application. New key associated data is stored within J6, J5A, and J5 format datasets as required to support this feature. Labels are not recorded by using the zero-key method in this mode.
- The reported Format Identifiers in Medium Sense reports non-encrypted J5A and J5 format identifier, unless all blocks on the medium are encrypted. Encryption format for a volume is determined at first write from BOT, and is enforced for all subsequent appended blocks. Labels are encrypted with zero key in this mode.

The TS1160, TS1155, TS1150, and TS1140 tape drives support the following encryption management techniques for open systems:

- Application-managed encryption (AME)
- Library-managed encryption (LME)
- System-managed encryption (SME) (for z/OS supported drives only)

Note: System-managed encryption (SME) is supported in the TS4500 for TS1140, and TS1150 drives only. This is provided for TS7700 IBM z Systems attachment only.

For more information about encryption, see Chapter 3, “Encryption” on page 173.

IBM Security Guardium Key Lifecycle Manager

IBM Security Guardium Key Lifecycle Manager (GKLM), former known as IBM Security Key Lifecycle Manager (SKLM), is the IBM strategic platform for the storage and delivery of encryption keys to encrypt storage endpoint devices. IBM GKLM can be used with the following tape drives:

- TS1160
- TS1155
- TS1150
- TS1140

Similar to the previous product, IBM Encryption Key Manager (EKM) and IBM Security Key Lifecycle Manager (SKLM), GKLM serves data keys to the tape drive.

It focuses on ease of use and provides a graphical user interface (GUI) to help with the installation and configuration of the key manager. It also allows for the creation and management of the key encrypting keys (certificates).

For more information about IBM Security Guardium Key Lifecycle Manager, see this IBM Documentation web page.

2.1.11 IBM Spectrum Archive (LTFS) support

IBM Spectrum Archive, a member of the IBM Spectrum Storage family, provides direct, intuitive, and graphical access to data that is stored in tape drives and libraries, including the latest LTO 9 tape technology. It incorporates the Linear Tape File System (LTFS) format standard for reading, writing, and exchanging descriptive metadata on formatted tape cartridges.

Spectrum Archive eliminates the need for extra tape management and software to access data. Spectrum Archive offers the following software solutions for managing your digital files with the LTFS format:

- IBM Spectrum Archive Single Drive Edition (SDE)

  This edition allows access to all data on a tape cartridge in a stand-alone drive that is attached directly to a server or through bridge box as easily as though it were on disk. IBM Spectrum Archive SDE uses the file system's format and operating system resources to graphically display the contents of a tape cartridge in the operating system's GUI format, typically a folder-tree structure.

  After they are mounted, the metadata of each cartridge is cached in server memory. Metadata operations, such as directory browsing and filename search, do not require tape movement. Similar to the use of a USB drive or memory stick, IBM Spectrum Archive SDE enables users to share data across platforms by dragging and dropping files.

- IBM Spectrum Archive Library Edition (LE)

  This edition provides multiple tapes in the LTFS format as an LTFS file system on a server. Each tape is a separate subdirectory in the LTFS file system. Files that are written to the LTFS file system must be written to one of the tape subdirectories.

  If the tape that represents the subdirectory is full, the write operation fails and the user must select a different subdirectory.

  Another aspect of IBM Spectrum Archive LE is that all read and write operations are done on the tape; no buffering or "housekeeping" operations, such as reclamation and reconciliation that must be conducted manually, are needed.

- IBM Spectrum Archive Enterprise Edition (EE)

  This edition extends the IBM Spectrum Scale file system name space to tapes in the LTFS format that are provided by an LTFS file system. The user interacts with the IBM Spectrum Scale file system, which serves as buffer for files that are migrated and recalled to the LTFS file system. For file migration jobs, it automatically selects a tape that provides sufficient capacity.

  IBM Spectrum Archive EE integrates with the IBM Spectrum Scale policy engine for migration and recalls. It also has built-in housekeeping operations for reclamation, reconciliation tape repair operations

IBM Spectrum Archive presents tape media as though it were a disk file system. IBM Spectrum Archive supports the IBM LTO Ultrium 9, 8, 7, 6, and 5, and IBM TS1160, TS1155, TS1150, and TS1140 tape drives.

Tape as a storage medium offers many benefits. Tape is reliable, portable, low-cost, low-power, and high-capacity. However, tape is not easy to use. It has no standard format, and data often cannot be used unless the data is copied to disk first.

With IBM Spectrum Archive, accessing data that is stored on an IBM tape cartridge became as intuitive as the use of a USB flash drive. With IBM Spectrum Archive, reading data on a tape cartridge is as easy as dragging and dropping a file. Users can run any application that is designed for disk files against tape data without concern that data is physically stored on tape.

IBM Spectrum Archive implements a true file system for tape. It also supports library automation, including the ability to find data on a tape in a library without mounting and searching tape volumes.

IBM Spectrum Archive Library Edition (LE) supports IBM tape automation and the single drive edition IBM Linear Tape File System. With IBM Spectrum Archive LE, you can create a single file system mount point for a logical library that is managed by a single instance of the software, which is running on a single server. In addition, it provides for caching of tape indexes, and for searching, querying, and displaying tapes' contents within an IBM tape library without the requirement to mount tape cartridges.

IBM Spectrum Archive provides the following features:

- Reduce complexity in data management and access time by enabling a self-describing tape with a simple file index
- Files that can be dragged to and dropped from the tape
- A simple one-time installation
- IBM Spectrum Archive operates the tape media by using two partitions (meta or index data and file content)
- Lower TCO by leveraging cost effective tape storage
- Seamless data access in continuous name space
- Automated, policy based movement from disk to tape
- Tape optimized recall to accelerate retrieves
- Standardized LTFS format facilitates data exchange
- Support for transparent tape encryption
- Data protection through multiple copies on tape
- Support for immutable files on WORM tapes
- Two-site replication by stretch cluster or AFM IW
- Media export and import for data sharing or off-site storage
- Media health check with TS4500
- Easy administration and management

2.2 IBM TS1160 tape drive

The IBM TS1160 tape drive (which is also referred to as the *3592 Model 60E*, *60F*, or *60S)* is an enhanced sixth-generation tape drive of the IBM 3592 tape family. The TS1160 tape drive provides higher levels of cartridge capacity than the TS1150 model EH8 and TS1155 models 55E and 55F. It is designed to provide an increased capacity on JE media types compared to its predecessors.

Note: The TS1160 model 60G drive is designed for installation in the TS3500.

A media-to-drive summary is shown in Figure 2-7 on page 97 for IBM Enterprise 3592 media types.

Note: TS1160 models require different TS4500 firmware levels:

- TS1160 model 60F requires 1.5 or higher
- TS1160 model 60E requires 1.6 or higher
- TS1160 model 60S requires 1.7.0.1 or higher

The TS1160 has the following key features, including those features that were introduced with the TS1155 and previous models:

- Digital speed matching
- Channel calibration
- High-resolution tape directory
- Recursive accumulating backhitch-less flush or nonvolatile caching (NVC)
- Backhitch-less backspace
- Virtual backhitch
- Read ahead
- Streaming Lossless Data Compression (SLDC)
- Capacity scaling
- Single FRU
- Error detection and reporting
- SARS
- Revised encryption support
- Dual-stage 32-head actuator
- Offboard data string searching
- Enhanced logic to report logical end of tape
- Added partitioning support
- Data Safe mode
- Enhanced Ethernet support
- New enhanced Barium Ferrite (BaFe) particle media types
- Dual port, 16 Gb FC attachment with failover support for FC drives
- Dual port, 10 or 25 Gb Ethernet ports for TS1160 model 3592 60E

- Dual port, 12 Gb SAS (Serial Attached SCSI) ports for TS1160 model 3592 60S
- Max Capacity mode logical end-of-tape (LEOT) support for up to 4% more capacity
- Partitioning that is supported by Spectrum Archive and IBM Linear Tape File System (LTFS)

2.2.1 TS1160 physical characteristics

The TS1160 drives have an identical form factor, and is plug-compatible with 3592 models. It maintains low power and improves power management. The drive power usage is 56 watts maximum operating power, as compared to 51 watts for the TS1140. Standby power is less than 23 watts.

The drives have a standby cooling management feature, which reduces fan speed when the drive is idle. It further reduces power and airborne debris contaminants. The fan operating mode is controlled by a single input signal that is called *full-speed mode* or *variable-speed mode*. In full-speed mode, the fan or blower runs at full speed. In variable-speed mode, the blower adjusts its speed based on the ambient temperature down to a minimum of about 50% of its full speed.

The speed of the fan is based on the following conditions:

- The drive code enables variable-speed mode under the following conditions:
  - The drive is unloaded and idle for 5 minutes.
  - The internal temperature is at least 3 degrees below the full speed required temperature limit.
- The drive code reverts to full-speed mode as soon as the following conditions are met:
  - A cartridge is placed in the loader or loaded.
  - The internal temperature of the drive rises above the full speed required temperature limit.

The internal temperature sensor is sampled at 5-minute intervals.

Internal hardware enhancements

These drives feature the following significant hardware enhancements over the previous models:

- New dataflow ASIC chip (Bara) in CU-32 technology
- Enables higher data rate of 400MBps)
- Enables longer C2 ECC code and iterative decoding (effective SNRa uplift)
- Increases HIB transfer rate to 800 MBps for FC-16/Ethernet attachment
- SPA write driver
- TMR read sensor technology of approximately 1um reader width
- High Bs writer to support writing of higher capacity JE tape
- TMR Shorting mitigators TBALL Readers
- New FC-16 Lancer G6 host target chip and PCIE bridge FPGA
- New Arrowhead Qlogic Ethernet chip with support for iSCSI/RoCE 10 Gb, 25 Gb
- Spring load rollers with new lubricant
- New JE media support
- Magnetic layer features new higher Hc perpendicularly oriented BaFe particle
- Thinner underlayer that supports longer tape.

The data-dependent, noise-predictive, maximum-likelihood (DD-NPML) detection scheme was developed at IBM Research®, Zurich, to enable the accurate detection of data errors.

2.2.2 Media

The TS1160 drives use the new higher Hc perpendicularly oriented BaFe particle media types.

This new media uses oriented perpendicular BaFe mag layer, which is an approximately 1.8 dB bbSNR improvement from JC/JD. The new media can be read/written up to 400 MBps native sustained data rate (up to 900 MBps at 3:1 compression ratio) in the new 32-channel Jag6, 6E, 5, 5E, 5A, and 5AE logical format, as listed in Table 2-8 on page 108.

For more information about TS1160 supported media types, see “3592 media cartridge” on page 94.

2.2.3 Capacity and performance

Capacity and performance were improved from the TS1160 tape drive for all media types, and for all formats that the drives reads or writes. These improvements require the TS1160 to format the tape or write from BOT.

Capacity improvement

The use of the 60F logical format offers the following capacity improvements on existing and new cartridges:

- IBM Enterprise Advanced Data media (JE and JV), which has a capacity of 20 TB
- IBM Enterprise Advanced Data media (JD and JZ), which has a capacity of 15 TB
- IBM Enterprise Advanced Data media (JC and JY), which has a capacity of 7 TB
- IBM Enterprise Economy Data media (JM), which has a capacity of 5 TB
- IBM Enterprise Economy Data media (JL), which has a capacity of 3 TB
- IBM Enterprise Economy Data media (JK), which has a capacity of 900 GB

Performance

The overall performance is improved over the previous model by the following improvements:

- Improved data rate and capacity
- Improved latency by reducing access time to data
- Increases HIB transfer rate to 1200 MBps
- Beginning of partition (BOP) caching
- Humidity sensor support
- Increased cartridge memory size and related functions
- Improved high-resolution tape directory (HRTD)
- New dataflow ASIC chip in CU-32 technology
- Extended copy support

Higher data rates and capacity

The 60E and 60F format data rates go up to 400 MBps maximum native, and to 900 MBps maximum compressed (at 256 K and greater block size). Table 2-8 on page 108 lists the capacity and performance characteristics for uncompressed data.

Table 2-8 Capacity and performance summary

Media	60F format capacity native data rate (minimum - maximum)	55F format capacity native data rate (minimum - maximum)	EH8 format capacity data native rate (minimum - maximum)
JE or JV	20 TB (18.12 TiB) 122 MBps - 407 MBps	N/A	N/A
JC or JY	7 TB (6.37 TiB) 99 MBps - 303 MBps	7 TB (6.37 TiB) 99 MBps - 303 MBps	7 TB (6.37 TiB) 99 MBps - 303 MBps
JD or JZ	15 TB (13.64 TiB) 112 MBps - 365 MBps	15 TB (13.64 TiB) 112 MBps - 365 MBps	10 TB (9.1 TiB) 112 MBps - 365 MBps
JM	5 TB (4.55 TiB) 122 MBps - 407 MBps	N/A	N/A
JL	3 TB (2.73 TiB) 112 MBps - 365 MBps	3 TB (2.73 TiB) 112 MBps - 365 MBps	2 TB (1.82 TiB) 112 MBps - 365 MBps
JK	900 GB (.82 TiB) 99 MBps - 303 MBps	900 GB (.82 TiB) 99 MBps - 303 MBps	900 GB (.82 TiB) 99 MBps - 303 MBps

Improved latency

These tape drives add the following features to improve latency by reducing access time to data:

- Improved locate and rewind speed profile for the new media types by using 12.4 meters (13.5 yards) per second (m/s) end-to-end versus 12.4 m/s profiled (JE, JD, JV, JZ, JM, and JL media only):
  - JE, JD, JZ, JM, and JL media feature a redesigned brake button for higher reliability, longer life, and higher locate speeds.
  - The improved profile represents a 9% speed improvement for a rewind/locate operation from EOT to BOT versus the previous profile, which partially compensates for the longer tape length of new media types.
- Load and thread times are reduced by approximately 33% from 15 seconds load/ready to 10 seconds load/ready. This reduction applies to both JC, JD, and JE media types.

This improvement is possible by operating motors at a higher operating speed for repeatable read (RR), loader, and threader motors.

Compression

The TS1160 drives feature the same history buffer usage in the compression core as TS1155. The history buffer is 16 KiB, which enables more efficient compression by increasing the history over which string matches can be applied. The new method increases the nominal compression ratio for the Calgary Corpus data standard from approximately 2.0 to 2.5.

> **Note:** The improved compression method is available only when processing the Jag-5/6 logical formats on the new JE/JV/JM media or unformatted JD/JZ/JL/JC/JY/JK media. When processing the Jag-4 format, an older compression method is used for compatibility.

As in previous models, the 3592 tape drive uses data compression that is known as *Streaming Lossless Data Compression Algorithm* (SLDC). This compression method is identical to the method that was used in previous models, except for the larger history buffer.

SLDC is an implementation of a Lempel-Ziv class 1 (LZ-1) data compression algorithm. SLDC also is an extension to the Adaptive Lossless Data Compression (ALDC) algorithm, which is used in leading industry tape products. Users of SLDC can expect to achieve the same, or better, data compression as users of ALDC.

A key difference between SLDC and previous lossless compression algorithms is that record boundaries and file marks are encoded as control symbols. The encoding of record boundaries and file marks as control symbols allows the compressed data stream to be separated into a serial stream of records and file marks by the decompression logic without requiring additional information, such as information from an attached header.

Beginning of partition caching

These drives implement beginning of partition (BOP) caching. In this implementation, after the initial set of tape blocks in a partition is read by the read-ahead function or an explicit command, the initial set of tape blocks remains in a special place in the cache data buffer (until an unmount or a partition change). Subsequent locate operations to BOP or read operations of these blocks complete quickly, without requiring completion of physical motion. BOP caching is supported in all partition modes.

This feature is automatic, cannot be disabled, and uses approximately 6 MB space (one data set) in the main data buffer.

Humidity sensor

The drives contains a humidity sensor and a temperature sensor. The humidity sensor provides the following functions:

- Humidity tracing in drive logs

  The drive logs humidity data in the tape map during read and write.

- Maximum humidity logging in cartridge memory

  The maximum humidity that is sensed during a cartridge mount is loaded in the cartridge memory.

- Humidity data is externalized in log pages and as with temperature data, humidity data can now be read through standardized SCSI Log pages by an initiator. However, environmental thresholds cannot be set.

Improved high-resolution tape directory

The TS1160 drive provides a higher-granularity directory to improve the accuracy of tape locate operations for new JE, JV, and JM media types. The granularity of wrap entries is unchanged from JD media. High-resolution tape directory (HRTD) resolution for JC media types is unchanged. HRTD directories are maintained separately for partitions.

They maintain a tape directory structure with a high granularity of information about the physical position of data blocks and file marks on the media. The longitudinal position (LPOS) longitudinal location information that is contained in the servo pattern is associated with and recorded with the host block information in the HRTD. This feature allows the 3592 to have fast and consistent nominal and average access times for locate operations.

Therefore, locate times are uniform and based on the position of the block or file mark on the tape independent of the uniformity of the block size or file mark distribution along the length of the tape.

The HRTD feature maintains an overall granularity of 64 directory entries per logical wrap. Consider the following points:

- JA media 570 m (623.36 yards) logical wrap results in a granularity of 8.9 meters (29.1 ft.).
- JB media 775 m (847.55 yards) logical wrap results in a granularity of 12.1 meters (39.6 ft.).
- JC media 842 m (920.8 yards) logical wrap results in a granularity of 13.2 meters (43.3 ft.).
- JD media 1032 m (1128.6 yards) logical wrap results in a granularity of 8.06 meters (26.4 ft.). Granularity is improved for the segmented or scaled formats with shorter logical wraps.
- JE media 1088 m (1190 yards) logical wrap results in a granularity of 8.06 meters (26.4 ft.). For the segmented or scaled formats with shorter logical wraps, granularity is improved.

The 3592 drive has many redundancy and recovery features that prevent the possibility of data loss in the loss of a directory. It also allows a rebuild of the directory under all circumstances. consider the following points:

- The HRTD table consists of information for each logical wrap. Each wrap area contains up to 64 entries. Each entry contains the LPOS, logical block, and file mark count information with access point and other internal information of interest.
- The entire HRTD table is stored in the housekeeping data set on tape. The entire HRTD structure is also written in the end-of-data (EOD) data set for the tape if the tape has a valid EOD. The HRTD entries are also distributed in accumulating sequential fashion into the Data Set Information Table of all user data sets as they are written on tape. Control structures, which define the validity of the HRTD and EOD information on the tape, are in the cartridge memory.
- If a valid HRTD cannot be recovered from the housekeeping data set, the HRTD might be rebuilt by using the EOD or distributed copies of HRTD information. The HRTD can also be rebuilt by reading the tape. Depending on the mechanism that must be used to rebuild the HRTD, this rebuild occurs quickly (seconds if the EOD copy can be used) or take longer (minutes if a full rebuild is required).
- The drive can read all data from a cartridge without any HRTD information, although locate times might be affected. However, the drive does not allow a write operation without a valid HRTD to guarantee the integrity and validity of the information on tape.

Main data buffer

These drives feature the same 2 GB main data buffer as the TS1155, which is twice the size of the 1 GB main buffer in the TS1140 drive. The extra buffer is used to improve overall performance, reduce backhitches, improve speed matching performance, and support BOP caching and other improvements.

External copy support

These drives support the *external copy* function, which offers the following advantages:

- The capability is similar to serverless copy in that it allows data to be copied from one drive to another drive with no transfer through the host at high data rates.
- Data can be an entire volume or a group of logical blocks.
- The hosting drive (TS1160 or any drive that supports the feature) can pull or push data to a second drive of any type (vendor-neutral and does not require feature support).
- The function works in a SAN environment, and it is supported on true switches (non-hubs).

SkipSync or Same Wrap Backhitchless Flush mode feature

These drives, as with previous models, implement a feature that is known as *same wrap backhitchless flush mode* (SWBF mode), which is also called the *SkipSync* feature. This feature is similar to previous models, plus the following enhancements:

- In default mode, SkipSync is enabled to use up to 1.5% capacity loss and uses spare capacity, so no impact occurs to the client capacity in the nominal Constant Capacity LEOT mode.
- SkipSync can be programmed through Mode page 0x30 to allow up to 33% capacity loss, which essentially enables SkipSync for all transactions.
- The performance (throughput) improves for operations or transaction sizes that use SkipSync because of the increased nominal data rates of the TS1160.

How SkipSync operates

When a sync command (`WFM 0`) or a Write File Mark (WFM) non-immediate command is received after a block or series of data blocks (referred to here as a *transaction*), the TS1160 drive does not perform a backhitch immediately after the synchronization or WFM completes. Instead, it continues to stream on the same wrap and write a Data Set Separator (DSS) pattern until enough data is received to record more data sets. SkipSync results in a significant performance improvement because of backhitch avoidance, but a reduction in the overall available capacity on the volume.

In default mode, SWBF mode (SkipSync) is entered after a flush is received under the following conditions:

- The received transaction size is greater than 204 MB compressed.
- The drive is not already in Recursive Accumulating Backhitchless Flush (RABF) mode.
- Enough excess capacity remains based on the current LPOS so that the drive predicts that it will still achieve the minimum capacity threshold that is selected. The minimum capacity threshold is 1.5% for the TS1160 default mode.

Virtual backhitch

These drives include the following key feature improvements:

- Virtual backhitch (transaction write with sync)
- Single wrap backhitchless flush (large transaction writes with sync)
- Backhitchless backspacing (American National Standards Institute (ANSI) file writes)

The TS1160 function uses Recursive Accumulating Backhitchless Flush (RABF) and the addition of a new same wrap backhitchless flush (SWBF) function that extends virtual backhitch effectiveness for large files.

Fast sync and skip performance for these tape drives are enhanced because of the better data rate performance over the TS1140.

For more information about these features, see “Virtual backhitch (nonvolatile caching)” on page 84.

Read ahead feature

On sequential reads, the tape drive automatically runs *read ahead* and fills the buffer with data sequentially beyond the target block.

These drives support advanced automatic read-ahead and read-space virtualization at improved access performance and 2x data buffer size. When the drive processes a command to locate or read a block, the drive automatically continues to stream down the tape and read ahead until the data buffer is full. This feature allows subsequent `Locate` or `Read` commands to be fulfilled from the data buffer at faster speeds, rather than requiring access to the tape.

With this unique function, the drive outperforms competitive drives, which stop and wait for the next command.

2.2.4 Emulation

The TS1160 supports drive emulation, but not emulation mode.

Emulation mode

The TS1150 does not support any emulation modes. Because the drive cannot write the TS1130, TS1120, or J1A logical format, it cannot fully emulate all format behaviors of a previous model 3592 drive.

Drive emulation

The TS1160 tape drive does not support emulation. The TS1160 can read and write in J5A and J5 format with compatible IBM 3592 tape cartridges.

The TS1160 tape drive can reformat any compatible J5 tape when it is writing from BOT and the TS1160 can reformat any J5 format tape. Table 2-9 on page 112 lists the available modes for TS1160, TS1150, and TS155.

Table 2-9 Drive emulation for TS1160

Drive mode setting	Formats read	Format that is used when the writing cartridge is at BOT	Format that is used when the writing cartridge is not at BOT	Model type that is reported to the host in response to the Inquiry command
EH8 J5 format	J4 J5	J5	J5 if format at J5 J4 if format at J4	E08
55E and 55F J5A format	J5 J5A	J5A	J5A if format at J5A J5 if format at J5	55E or 55F
60E, 60F, and 60S J6 format	J5 J5A J6	J6	J6 if format at J6 J5A if format at J5A J5 if format at J5	60E, 60F, or 60S

The TS1155 and TS1150 can reformat a compatible tape written in J6 format, but cannot read in this format.

2.2.5 Data safe mode

The TS1160 supports *data safe mode*. This mode is controlled by the application and prevents inadvertent overwrite. Data safe mode treats the tape volume that is mounted as a WORM drive and prevents inadvertent overwrite. This mode is set by the application or host system.

2.2.6 Upgrade considerations

A drive-field *Miscellaneous Equipment Specification* (MES) conversion feature is available for a 3592-EH7 model to 3592-EH8 model conversion to a TS1150.

> **Important:** If you choose this MES to replace the TS1140 drive, only the drive changes. The canister remains the same. The serial number of the original drive is written by the library to the vital product data (VPD) of the replacement drive. The MES is valid for both the TS4500 tape library and a rack-mounted drive.

TS1150 Field MES support

The following drive MES conversions are supported:

- 3592 EH8 drive > 3592 60F model upgrade
- 3592 55F drive > 3592 60F model upgrade
- 3592 60G drive > 3592 60F model upgrade
- 3592 55G drive > 3592 60G model upgrade
- 3592 E08 drive > 3592 60G model conversion

> **Important:** You must load the IBM Assembly and Deploy Tools (IADT)-capable microcode for a TS1140 drive before conversion to the EH8 model. Without it, communication to the library is not possible. The IADT-capable microcode can be obtained from IBM Fix Central. Only one MES model upgrade is supported in the life of the drive.

2.2.7 Firmware updates

No changes were made to the firmware update mechanisms for the TS1160 as compared to previous TS1100 tape drives. Consider the following points:

- The TS1160 continues to support concurrent Licensed Internal Code (LIC) load with deferred activation.
- The TS1160 have LIC images that are unique from previous models.
- The TS1160 models 60E, 60F, and 60S have each a unique firmware ID.

The firmware for the TS1160 drives can be updated by using one of the following methods, depending on where the drive is installed:

- Through the TS4500 management GUI
- Through the host attachment by using the write buffer command or IBM TotalStorage Tape Diagnostic Tool (ITDT)

2.2.8 RAS

The RAS features are improved or maintained relative to the TS1160. They are similar to their predecessor models (the TS1160 is a single FRU), which are hot-pluggable without a maintenance window and support nondisruptive code loading. As with the TS1140, fan speed management and unique device microcode file management are available through a LOAD ID.

The end of life usage alert for media activates on full-file pass usage. The Nearing Media Life alert occurs at 19,900 mounts or 295 full-file passes. The Media Life alert for JD or JL media use within a TS1150 drive is now rated for 20 M motion meters as opposed to 300 Full-File Passes (FFPs). The Media Life alert for JE, JM, or JV media that is used within a TS1160 drive is now rated for 100 Full-File Passes (FFPs).

2.3 IBM TS1155 and TS1150 tape drive

The IBM TS1155 tape drive (which is also referred to as the *3592 Model 55G, 55E, or 55F)* is an enhanced fifth-generation tape drive of the IBM 3592 tape family. The TS1155 tape drive provides higher levels of cartridge capacity than the TS1150 Model E08 (EH8). It is designed to provide an increased capacity of 50% on JD media types compared with its predecessors.

Note: The TS1155 model 55G and TS1150 E08 are drives designed for installation in the TS3500.

The primary difference of the TS1155 (55E, 55F) from the base TS1150 drive (EH8) is that the capacity is increased on JD media types.

The TS1155 model 55E replaces the FC host attachment by dual 10 GB Ethernet (RoCE v2) ports, which is optimized for cloud-based and large, open-compute environments.

The TS1155 model 55F and TS1150 model EH8 tape drives have a dual-port, 8 Gbps Fibre Channel interface for Fibre Channel attachment to host systems, or a switched fabric environment.

The TS1155 Tape drive is not compatible with IBM TS7700 or Enterprise Tape Control Unit environments.

The TS1155 and TS1150 have the following key features, including those features that were introduced with the 3592-J1A, 3592-E05, 3592-E06, and 3592-E07 (EH7):

- Digital speed matching
- Channel calibration
- High-resolution tape directory
- Recursive accumulating backhitch-less flush or nonvolatile caching (NVC)
- Backhitch-less backspace
- Virtual backhitch
- Read ahead
- Streaming Lossless Data Compression (SLDC)
- Capacity scaling
- Single FRU
- Error detection and reporting
- SARS
- Revised encryption support
- Dual-stage 32-head actuator

- Offboard data string searching
- Enhanced logic to report logical end of tape
- Added partitioning support
- Data Safe mode
- Enhanced Ethernet support
- New enhanced Barium Ferrite (BaFe) particle media types
- Dual-port, 8 Gb FC attachment with failover support for FC drives
- Dual port, 10 Gb Ethernet ports for TS1155 55E
- Max Capacity mode logical end-of-tape (LEOT) support for up to 4% more capacity
- Partitioning that is supported by Spectrum Archive and IBM LTFS

2.3.1 TS1155 and TS1150 physical characteristics

The TS1155 and TS1150 drives have an identical form factor, and it is plug-compatible with existing 3592 models. It maintains low power and improves power management. The maximum continuous operating power was decreased by 5 watts from the TS1140. The drive power usage is 46 watts maximum operating power, as compared to 51 watts for the TS1140. Standby power is fewer than 23 watts.

They have a standby cooling management feature, which reduces the fan speed when the drive is idle to further reduce power and reduce airborne debris contaminants. The fan operating mode is controlled by a single input signal that is called *full-speed mode* or *variable-speed mode*. In full-speed mode, the fan or blower runs at full speed. In variable-speed mode, the blower adjusts its speed based on the ambient temperature down to a minimum of about 50% of its full speed.

The speed of the fan is based on the following conditions:

- The drive code enables variable-speed mode under the following conditions:
  - The drive is unloaded and idle for 5 minutes.
  - The internal temperature is at least 3 degrees below the full speed required temperature limit.
- The drive code reverts to full-speed mode as soon as the following conditions are met:
  - A cartridge is placed in the loader or loaded.
  - The internal temperature of the drive rises above the full speed required temperature limit.

The internal temperature sensor is sampled at 5-minute intervals.

Internal hardware enhancements

These drives feature the following significant hardware enhancements over the previous models:

- A 32 channel-enhanced error correction code (ECC) recording format.
- Enhanced JD-type media servo pattern.
- Flangeless rollers, which are designed to minimize tape edge damage and debris buildup by the elimination of the roller flanges.
- Tunnel Magneto Resistive (TMR head) on the TS1155 and on new version of the TS1150 whereas older generation of TS1150 use a Giant Magneto Resistive Heads (GMR head). These advanced heads are designed to reduce friction with advanced head coating to prevent corrosion and to extend head and tape cartridge life.

- Skew Actuator, which allows dynamic skew adjustment of the head to keep the head perpendicular to the tape.

The data-dependent, noise-predictive, maximum-likelihood (DD-NPML) detection scheme was developed at IBM Research, Zurich, to enable the accurate detection of data errors.

2.3.2 Media

The TS1155 and TS1150 drives use the following enhanced Barium Ferrite (BaFe) second-generation particle media types. The new media can be read/written up to 360 MBps native sustained data rate (up to 700 MBps at 3:1 compression ratio) in the new 32-channel Jag 5, 5E, 5A, and 5AE logical format, as listed in Table 2-10.

Table 2-10 Media type and read/write format compatibility

Media type	TS1155				TS1150				
	Logical Format read/write				Logical Format read/write				z/OS Media Type
JD	J5A	J5A-E	J5	J5-E	J5	J5-E			Media 14
JZ	J5A	J5A-E	J5	J5-E	J5	J5-E			Media 15
JL	J5A	J5A-E	J5	J5-E	J5	J5-E			Media 16
JC	J5A	J5A-E	J5	J5-E	J5	J5-E	J4	J4-E	Media 11
JY	J5A	J5A-E	J5	J5-E	J5	J5-E	J4	J4-E	Media 12
JK	J5A	J5A-E	J5	J5-E	J5	J5-E	J4	J4-E	Media 13

Important: The TS1155 and TS1150 are not compatible with several older 3592 cartridge media types: JA, JB, JW, JJ, JR, and JX media types J3, J2, and J1. (MEDIA5, MEDIA6, MEDIA7, MEDIA8, MEDIA9, and MEDIA10).

TS1155 is not supported on z/OS.

These drives improve capacity and performance by writing and reading J5 and J5A logical format and by using a new 32-channel enhanced ECC recording format with a higher track density and higher linear density on the same media types.

The appropriate microcode levels that are available for TS1150 and TS1140 and must be installed that enable the recognition of the J5A and J5 format and allow reuse of the media in the older formats. Therefore, a model J5A drive can reformat media that was written in the older format, and write on it in the appropriate format.

Important: This design supports a common scratch pool by media type, regardless of the last written format or allocation target drive.

2.3.3 Capacity and performance

Capacity and performance were improved from the TS1155 and TS1150 tape drive for all media types, and for all formats that the drives reads or writes.

Capacity improvement

The use of the 3592-55E and 55F logical format offers the following capacity improvements on existing and new cartridges:

- IBM Enterprise Advanced Data media (JZ and JD), which is a capacity of 15 TB
- IBM Enterprise Advanced Data media (JC and JY), which is a capacity of 7 TB
- IBM Enterprise Economy Data media (JZ), which is a capacity of 3 TB
- IBM Enterprise Economy Data media (JK), which is a capacity of 900 GB

The use of the 3592-EH8 logical format offers the following capacity on existing and new cartridges:

- IBM Enterprise Advanced Data media (JZ and JD), which is a capacity of 10 TB
- IBM Enterprise Advanced Data media (JC and JY), which is a capacity of 7 TB
- IBM Enterprise Economy Data media (JZ), which is a capacity of 2 TB
- IBM Enterprise Economy Data media (JK), which is a capacity of 900 GB

Performance improvement

The overall performance is increased over the previous model by various improvements:

- Improved data rate and capacity
- Improved latency by reducing access time to data
- Improved data compression
- Beginning of partition (BOP) caching
- Humidity sensor support
- Increased cartridge memory size and related functions
- Improved high-resolution tape directory (HRTD)
- Larger main data buffer
- Extended copy support

Higher data rates and capacity

The following format data rates are available (at 256 K and greater block size):

- The 55E format data rates go up to 360 MBps maximum native, and to 600 MBps maximum compressed.
- The EH8 and 55F format data rates go up to 360 MBps maximum native, and to 700 MBps maximum compressed.
- The EH7 format data rates go up to 250 MBps maximum native, and to 700 MBps maximum compressed.

The capacity and performance characteristics for uncompressed data are listed in Table 2-11.

Table 2-11 Capacity and performance summary

Media	55E or 55F format capacity native data rate (minimum - maximum)	EH8 format capacity data native rate (minimum - maximum)	EH7 format capacity native data rate (minimum - maximum)
JC or JY	7 TB (6.37 TiB) 99 MBps - 303 MBps	7 TB (6.37 TiB) 99 MBps - 303 MBps	4 TB (3.6 TiB) 90 MBps - 251 MBps
JD or JZ	15 TB (13.64 TiB) 112 MBps - 365 MBps	10 TB (9.1 TiB) 112 MBps - 365 MBps	N/A
JK	900 GB (.82 TiB) 99 MBps - 303 MBps	900 GB (.82 TiB) 99 MBps - 303 MBps	500 GB (.48 TiB) 90 MBps - 251 MBps
JL	3 TB (2.73 TiB) 112 MBps - 365 MBps	2 TB (1.82 TiB) 112 MBps - 365 MBps	N/A

Improved latency

These tape drives adds features to improve latency by reducing access time to data:

- Improved locate and rewind speed profile for the new media types by using 12.4 meters (13.5 yards) per second (m/s) end-to-end versus 12.4 m/s profiled (JD, JZ, and JL media only):
  - JD, JZ, and JL media feature a redesigned brake button for higher reliability, longer life, and higher locate speeds.
  - The improved profile represents a 9% speed improvement for a rewind/locate operation from EOT to BOT versus the previous profile, which partially compensates for the longer tape length of the new media types.
- Load and thread times are reduced by approximately 33% from 15 seconds load/ready to 10 seconds load/ready. This reduction applies to both JC and JD media types.

This improvement is possible by operating the motors at a higher operating speed for repeatable read (RR), loader, and threader motors.

Improved data compression

The TS1155 and TS1150 drives feature a larger history buffer usage in the compression core. This change increases the history buffer 1 - 16 KiB, which enables more efficient compression by increasing the history over which string matches can be applied. The new method can increase the nominal compression ratio for the Calgary Corpus data standard from approximately 2.0 to 3.1.

> **Note:** For the TS1150, the improved compression method is only available when you process the EH8 logical formats on the new JD, JZ, and JL media or on unformatted JC, JY, or JK media. When you process the EH7 format, the historical compression method is used for compatibility. TS1155 cannot write in EH7 format.

As in previous models, the 3592 tape drive uses the data compression that is known as *Streaming Lossless Data Compression Algorithm* (SLDC). This compression method is identical to the method that was used in previous models, except for the larger history buffer. SLDC is an implementation of a Lempel-Ziv class 1 (LZ-1) data compression algorithm. SLDC is an extension to the Adaptive Lossless Data Compression (ALDC) algorithm, which is used in leading industry tape products. Users of SLDC can expect to achieve the same, or better, data compression as users of ALDC.

A key difference between SLDC and previous lossless compression algorithms is that record boundaries and file marks are encoded as control symbols. The encoding of record boundaries and file marks as control symbols allows the compressed data stream to be separated into a serial stream of records and file marks by the decompression logic without requiring additional information, such as information from an attached header.

Beginning of partition caching

These drives implement beginning of partition (BOP) caching. In this implementation, after the initial set of tape blocks in a partition is read, either by the read-ahead function or an explicit command, the initial set of tape blocks remains in a special place in the cache data buffer (until an unmount or a partition change). Subsequent locate operations to BOP or read operations of these blocks complete quickly, without requiring completion of physical motion. BOP caching is supported in all partition modes.

This feature is automatic, cannot be disabled, and uses approximately 6 MB space (one data set) in the main data buffer.

Humidity sensor

The drives contains a humidity sensor and a temperature sensor. The humidity sensor provides the following functions:

- Humidity tracing in drive logs

  The drive logs humidity data in the tape map during read and write.

- Maximum humidity logging in cartridge memory

  The maximum humidity that is sensed during a cartridge mount is loaded in the cartridge memory.

- Humidity data is externalized in log pages and like temperature data, humidity data can now be read through standardized SCSI Log pages by an initiator. However, environmental thresholds cannot be set.

Increased cartridge memory size and related functions

The new JD, JZ, and JL media types contain 16 KB cartridge memory (CM), which increased from the 8 KB that was contained in JC and JB media types. The CM contains a larger medium auxiliary memory (MAM) area, which is available to the application.

Improved high-resolution tape directory

The TS1155 and TS1150 drive provides a higher-granularity directory to improve the accuracy of tape locate operations for the new JD, JZ, and JL media types. The granularity of wrap entries is increased from 64/wrap to 128/wrap for the new media types. High-resolution tape directory (HRTD) resolution for JC media types is unchanged. HRTD directories are maintained separately for partitions.

They maintain a tape directory structure with a high granularity of information about the physical position of data blocks and file marks on the media. The longitudinal position (LPOS) longitudinal location information that is contained in the servo pattern is associated with and recorded with the host block information in the HRTD. This feature allows the 3592 to have fast and consistent nominal and average access times for locate operations.

Therefore, locate times are uniform and based on the position of the block or file mark on the tape independently of the uniformity of the block size or file mark distribution along the length of the tape.

The HRTD feature maintains an overall granularity of 64 directory entries per logical wrap. Consider the following points:

- JA media 570 m (623.36 yards) logical wrap results in a granularity of 8.9 meters (29.1 ft.).
- JB media 775 m (847.55 yards) logical wrap results in a granularity of 12.1 meters (39.6 ft.).
- JC media 842 m (920.8 yards) logical wrap results in a granularity of 13.2 meters (43.3 ft.).
- JD media 1032 m (1128.6 yards) logical wrap results in a granularity of 8.06 meters (26.4 ft.). Granularity is improved for the segmented or scaled formats with shorter logical wraps.

The 3592 drive has many redundancy and recovery features that prevent the possibility of data loss in the loss of a directory and allow a rebuild of the directory under all circumstances. Consider the following points:

- The HRTD table consists of information for each logical wrap. Each wrap area contains up to 64 entries. Each entry contains the LPOS, logical block, and file mark count information with access point and other internal information of interest.

- The entire HRTD table is stored in the housekeeping data set on tape. The entire HRTD structure is also written in the end-of-data (EOD) data set for the tape if the tape has a valid EOD. The HRTD entries are also distributed in accumulating sequential fashion into the Data Set Information Table of all user data sets as they are written on tape. Control structures, which define the validity of the HRTD and EOD information on the tape, are in the cartridge memory.
- If a valid HRTD cannot be recovered from the housekeeping data set, the HRTD might be rebuilt by using the EOD or distributed copies of HRTD information. The HRTD can also be rebuilt by reading the tape. Depending on the mechanism that must be used to rebuild the HRTD, this rebuild can occur quickly (seconds if the EOD copy can be used) or take longer (minutes if a full rebuild is required).
- The drive can read all data from a cartridge without any HRTD information, although locate times might be affected. However, the drive does not allow a write operation without a valid HRTD to guarantee the integrity and validity of the information on tape.

Larger main data buffer

These drives feature a 2 GB main data buffer, which is twice the size of the 1 GB main buffer in the TS1140 drive. The additional buffer is used to improve overall performance, reduce backhitches, improve speed matching performance, and support BOP caching and other improvements.

External copy support

These drives support the *external copy* function, which offers these advantages:

- The capability is similar to serverless copy in that it allows data to be copied from one drive to another drive with no transfer through the host at high data rates.
- Data can be an entire volume or a group of logical blocks.
- The hosting drive (TS1155 or any drive that supports the feature) can pull or push data to a second drive of any type (vendor-neutral and does not require feature support).
- The function works in a SAN environment, and it is supported on true switches (non-hubs).

SkipSync or Same Wrap Backhitchless Flush mode feature

These drives, like previous models, implement a feature that is known as same wrap backhitchless flush mode (SWBF mode), which is also called the *SkipSync* feature. This feature is similar to previous models, plus the following enhancements:

- In default mode, SkipSync is enabled to use up to 1.5% capacity loss and uses spare capacity, so no impact occurs to client capacity in the nominal Constant Capacity LEOT mode.
- SkipSync can be programmed through Mode page 0x30 to allow up to 33% capacity loss, which essentially enables SkipSync for all transactions.
- The performance (throughput) improves for operations or transaction sizes that use SkipSync because of the increased nominal data rates of the TS1150.

How SkipSync operates

When a sync command (`WFM 0`) or a Write File Mark (WFM) non-immediate command is received after a block or series of data blocks (referred to here as a *transaction*), the TS1150 drive does not perform a backhitch immediately after the synchronization or WFM completes. Instead, it continues to stream on the same wrap and write a Data Set Separator (DSS) pattern until enough data is received to record more data sets.

SkipSync results in a significant performance improvement due to backhitch avoidance but a reduction in the overall available capacity on the volume.

In default mode, SWBF mode (SkipSync) is entered after a flush is received under these conditions:

- The received transaction size is greater than 204 MB compressed.
- The drive is not already in Recursive Accumulating Backhitchless Flush (RABF) mode.
- Enough excess capacity remains based on the current LPOS so that the drive predicts that it still achieves the minimum capacity threshold that is selected. The minimum capacity threshold is 1.5% for the TS1150 default mode.

Virtual backhitch

These drives include the following key feature improvements:

- Virtual backhitch (transaction write with sync)
- Single wrap backhitchless flush (large transaction writes with sync)
- Backhitchless backspacing (American National Standards Institute (ANSI) file writes)

The TS1150 and TS1155 function is improved with Recursive Accumulating Backhitchless Flush (RABF) and the addition of a new same wrap backhitchless flush (SWBF) function that extends virtual backhitch effectiveness for large files.

Fast sync and skip performance for these tape drives are enhanced because of the better data rate performance over the TS1140.

For more information about these features, see “Virtual backhitch (nonvolatile caching)” on page 84.

Read ahead feature

On sequential reads, the tape drive automatically runs *read ahead* and fills the buffer with data sequentially beyond the target block.

These drives support advanced automatic read-ahead and read-space virtualization at improved access performance and 2x data buffer size. When the drive processes a command to locate or read a block, the drive automatically continues to stream down the tape and to read ahead until the data buffer is full. This feature allows subsequent `Locate` or `Read` commands to be fulfilled from the data buffer at faster speeds, rather than requiring access to the tape.

With this unique function, the drive outperforms competitive drives, which stop and wait for the next command.

2.3.4 Emulation

The TS1150 supports drive emulation, but not emulation mode.

Emulation mode

The TS1150 does not support any emulation modes. Because the drive cannot write the TS1130, TS1120, or J1A logical format, it cannot fully emulate all format behaviors of a previous model 3592 drive.

Drive emulation

The TS1155 and the TS1150 tape drive do not support emulation, The TS1155 can read and write in J5A and J5 format with compatible IBM 3592 tape cartridges. The TS1150 can read and write in J5 and J4 format with compatible IBM 3592 tape cartridges.

The TS1155 tape drive can reformat any compatible J5 tape when it is writing from BOT and the TS1150 can reformat any J4 format tape. Table 2-12 lists the available modes for TS1150 and TS155.

Table 2-12 Drive emulation for TS1150

Drive mode setting	Formats read	Format that is used when the writing cartridge is at BOT	Format that is used when the writing cartridge is not at BOT	Model type that is reported to host in response to the Inquiry command
EH8 J5 format	J4 J5	J5	J5 if format at J5 J4 if format at J4	E08
55E and 55F J5A format	J5A J5	J5A	J5A if format at J5A J5 if format at J5	55E or 55F

The TS1150 can reformat a tape written in J5A format, but cannot read in this format.

2.3.5 Data safe mode

The TS1155 and TS1150 supports *data safe mode*. This mode is controlled by the application and prevents inadvertent overwrite. Data safe mode treats the tape volume that is mounted as a WORM drive and prevents inadvertent overwrite. This mode is set by the application or host system.

2.3.6 Upgrade considerations

A drive-field *Miscellaneous Equipment Specification* (MES) is available for upgrading the 3592-EH7 model to a 3592-EH8 model.

Important: If you choose this MES to replace the TS1140 drive, only the drive changes. The canister remains the same. The serial number of the original drive is written by the library to the vital product data (VPD) of the replacement drive. The MES is valid for both the TS4500 tape library and a rack-mounted drive.

TS1150 Field MES support

The following drive MES conversions are supported:

- 3592 EH8 drive > 3592 55F model upgrade
- 3592 E07 drive > 3592 E08 model upgrade
- 3592 E07 drive > 3592 EH7 model upgrade
- 3592 EH7 drive > 3592 EH8 model upgrade
- 3592 E08 drive > 3592 EH8 model conversion

Important: You must load the IBM Assembly and Deploy Tools (IADT)-capable microcode for a TS1140 drive before the conversion to the EH8 model. Without it, communication to the library is not possible. The IADT-capable microcode can be obtained from IBM Fix Central. Only one MES model upgrade is supported in the life of the drive.

2.3.7 Firmware updates

No changes were made to the firmware update mechanisms for the TS1155 and TS1150 as compared to TS1140 tape drives:

- The TS1155 and TS1150 continues to support concurrent Licensed Internal Code (LIC) load with deferred activation.
- The TS1155 and TS1150 has a single LIC image that is unique from previous models.

Unique LIC is required for the model EH7 drives because the LOAD ID differs from the LOAD ID that is required for previous versions of 3592. The firmware for the 3592 55E, 55F, and EH8 drives can be updated by using one of the following methods, depending on where the drive is installed:

- Through the TS4500 management GUI
- Through the host attachment by using the write buffer command or IBM TotalStorage Tape Diagnostic Tool (ITDT)

2.3.8 RAS

The RAS features are improved or maintained relative to the TS1140. Similar to their predecessor models, the TS1155 and TS1150 are single FRU, which are hot-pluggable without a maintenance window and support nondisruptive code loading. As with the TS1140, fan speed management and unique device microcode file management are available through a LOAD ID.

The end of life usage alert for media activates on full-file pass usage. The Nearing Media Life alert occurs at 19,900 mounts or 295 full-file passes. The Media Life alert for JD or JL media use within a TS1150 drive is now rated for 20 M motion meters as opposed to 300 Full-File Passes (FFPs).

2.3.9 Improved media Statistical Analysis and Reporting System

The TS1155 and TS1150 supports Statistical Analysis and Reporting System (SARS) in a similar manner to previous drive models.

The tape drive uses SARS to help isolate failures between media and hardware. SARS uses the cartridge performance history (which is saved in the cartridge memory (CM) module) and the drive performance history (which is kept in the drive flash electrically erasable programmable read only memory (EEPROM)) to determine the likely cause of failure.

SARS can cause the drive to request a cleaning tape to mark the media as degraded, and to indicate that the hardware is degraded.

SARS information is reported through the TapeAlert flags and through media information messages (MIMs) or service information messages (SIMs).

The 3592 drive maintains a history of the last 100 mounts for both Volume Statistical Analysis and Reporting System (VSARS) and Hardware Statistical Analysis and Reporting System (HSARS).

Note: Media SARS information is preserved when media is reformatted.

The TS1150 implements an enhanced SARS function that is known as client-centric SARS (ccSARS). This function improves the overall amount of information that is maintained, and the presentation means to the client in concert with the automation system.

The media SARS function for the drives includes the following actions:

- Tape alerts are generated when media passes usage life, as determined by full-file passes, meters of tape that were processed, or the write pass count, and the total number of mounts (which was already supported).
- A media SARS summary is maintained in the cartridge memory in a manner where it can be rebuilt on tape if the SARS records on tape cannot be read and must be reinitialized. This cartridge memory copy is also readable on an earlier level TS1140 drive to preserve SARS information between logical format conversions.

2.3.10 Encryption

The TS1155 and TS1150 tape drives are encryption-capable. Like the TS1140, you do not need to enable the drive explicitly.

Encryption support includes the following enhancements:

- LME, AME, and SME (for TS7700).
- T10 default method support.
- Continued encrypted data key (EEDK) wrapped key support in LME and SME.
- Enhanced protocol support for Internet Printing Protocol (IPP), which can be configured for security, and Java Platform, Enterprise Edition 2 historical mode (as used by IBM Spectrum Protect), T10 default method, Security Protocol IN (SPIN), and Security Protocol OUT (SPOUT).
- Enhanced drive cryptographic upgrades to change the default authentication means from Secure Hash Algorithm-1 (SHA-1) to SHA-2 when you use IBM Security Guardium Key Lifecycle Manager.
- T10 standards-based encryption control on a logical block basis (not tied to format identifier) and writes encrypted data and clear data to the same tape cartridge.

2.4 IBM TS1140 tape drive (Model 3592 EH7)

The TS1140 tape drive offers a dual-port Fibre Channel host attachment interface. This feature provides flexibility in open systems environments because the drives can directly attach to open systems servers with Fibre Channel attachments. All TS1140 tape drives are encryption-capable.

> **Note:** Drive Model 3592 E08 and E07 or earlier drives' models are not supported by the HD2 frames of the TS4500 tape library.

The TS1140 Model EH7 tape drive has a dual-port 8 Gbps Fibre Channel interface for Fibre Channel attachment to host systems, or a switched fabric environment.

The TS1140 offers the following key features, including those features that were introduced with the 3592-J1A, 3592-E05, and 3592 E07:

- Digital speed matching
- Channel calibration
- High-resolution tape directory
- Recursive accumulating backhitchless flush or nonvolatile caching (NVC)
- Backhitchless backspace
- Streaming lossless data compression (SLDC) algorithm
- Capacity scaling
- Single field-replaceable unit (FRU)
- Error detection and reporting
- Statistical Analysis and Reporting System (SARS)
- Revised encryption support
- Dual-stage 32-head actuator
- Offboard data string searching
- Enhanced logic to report logical end of tape
- Added partitioning support
- End-to-end logical block protection support
- Data safe mode
- Enhanced Ethernet support
- New enhanced Barium Ferrite (BaFe) particle media types
- Eight Gbps Fibre Channel dual port interface

2.4.1 3592 media

In this section, we present more information about the TS1140 media.

Data cartridge

The TS1140 uses enhanced Barium Ferrite (BaFe) particle media types. The media can be read/written up to 250 MBps native sustained data rate (up to 650 MBps at 3:1 compression) in the 32-channel Generation-4 logical format. The following types of tape cartridges are available:

- IBM Enterprise Advanced Tape Cartridge (JC) - MEDIA11
- IBM Enterprise Advanced WORM Tape Cartridge (JY) - MEDIA12
- IBM Enterprise Advanced Economy Tape Cartridge (JK) - MEDIA13

The following TS1140 tape drive capacity and performance improvements are provided on the existing 3592 media:

- The TS1140 tape drive reuses certain TS1130 and TS1120 supported media types:
  - IBM 3592 Extended Tape Cartridge (JB) - MEDIA9
  - IBM 3592 Extended WORM Tape Cartridge (JX) - MEDIA10

  Note: Media types JA, JW, and JJ are not supported by the TS4500 tape library.

- The TS1140 improves capacity and performance by writing and reading the EH7 logical format by using a new 32-channel enhanced error correction code (ECC) recording format with a higher track density and higher linear density on the same media types.

  Design feature: This design supports a common scratch pool by media type regardless of the last written format or allocation target drive.

The 3592 tape drives can reuse different types of tape and multiple densities (logical formats) across various drive generations. Enterprise format 4 (EFMT4) records 2,560 tracks on 32 channels. These logical formats can be divided into multiple subformat options, such as segmentation and capacity scaling.

2.4.2 Capacity and performance

Capacity and performance are improved compared to the IBM TS1130 tape drive for all media types and for all formats that the TS1140 reads or writes.

Capacity improvement

The use of the 3592 EH7 logical format offers native capacity improvement on existing cartridges and new cartridges:

- IBM Enterprise Extended data (JB) or WORM (JX) media, a 60% capacity uplift from 1 TB to 1.6 TB (by using a 3:1 compression ratio 4.8 TB)
- IBM Enterprise Advanced data (JC) or WORM (JY) media, a capacity of 4.0 TB (by using a 3:1 compression ratio 12 TB)
- IBM Enterprise Economy Data media, JK media, a capacity of 500 GB (by using a 3:1 compression ratio 1.5 TB)

Performance improvement

The overall performance is increased by various improvements:

- Improved data rate
- Larger 1 GB main data buffer
- Better backhitching
- Improved speed with digital speed matching
- Enhanced read-ahead buffer management
- High access performance for locate or search
- Improved communication links with dual 8 Gbps Fibre Channel ports
- SkipSync and FastSync write performance accelerators
- New 32-channel enhanced ECC recording format

- IBM Spectrum Archive and IBM LTFS Single Drive Edition, Library Edition, and Enterprise Edition support

Note: *Backhitching* is the condition that occurs when a data cartridge stops, reverses, and restarts motion. A *backhitch* is the result of a mismatch between the data rates of the connected server and the tape drive.

Higher data rates

Performance is improved from the TS1130 up to 64% in TS1140 mode, 50% in TS1130 mode for read and writes, and 50% in TS1120 mode for reads only:

- The EH7 format data rates go up to 250 MBps maximum native and to 650 MBps maximum compressed.
- The E06 format data rates go up to 200 MBps maximum native and to 650 MBps maximum compressed.

The capacity and performance characteristics for native data are listed in Table 2-13.

Table 2-13 Capacity and performance summary

Media	EH7 format capacity data rate (minimum - maximum)	E06 format capacity data rate (minimum - maximum)
JB and JX	1.6 TB (1.46 TiB) 80 MBps - 200 MBps	1 TB (.9 TiB) 50 MBps - 160 MBps
JC and JY	4 TB (3.6 TiB) 90 MBps - 250 MBps	N/A
JK	500 GB (.48 TiB) 60 MBps - 250 MBps	N/A

Buffer

The TS1140 EH7 tape drive has a 1 GB internal data buffer. In addition to enabling higher performance characteristics, the data buffer is designed to support the “read-ahead” capability of compressed data from tape. The data buffer provides high-performance random skip forward sequential (short hop) locates, which are common in database search and tape software recycle operations.

This buffer improves the drive agility, file access, and small file handling. Furthermore, the buffer reduces backhitches for all workloads and improves overall read/write performance.

Offboard data string searching

The TS1140 EH7 tape drive can search the data content of host records for string matches offboard from the host server. The tape drive can perform this search at the maximum data rate (250 MBps native). It takes much longer for a host server to read the data, buffer the data to disk, and then parse the actual data stream with host software routines.

Enhanced logic to report logical end-of-tape (LEOT)

LEOT is now reported, based on a combination of capacity-based and position-based LEOT indicators. The TS1140 EH7 monitors the total accumulated number of physical tape data sets that is written to the volume and reports LEOT based on this capacity-based LEOT value. LOET allows tape copies to complete without overflow more often.

High-resolution tape directory

The TS1140 EH7 tape drive maintains a tape directory structure with a high granularity of information about the physical position of data blocks on the media. This high-resolution tape directory (HRTD), plus the increased search speed, improves the TS1140 EH7 nominal and average access times for locate operations.

Channel calibration and dynamic adaptive equalization

To gain optimum performance, *channel calibration* allows the drive to automatically customize each read and write data channel. The customization compensates for variations in the recording channel transfer function, for media characteristics, and for read and write head characteristics. Initial calibration settings are calculated and stored at the time of manufacture. For optimum error rate performance, the TS1140 tape drives also use dynamic adaptive equalization hardware on an ongoing basis to adjust the read equalization response.

Virtual backhitch

The TS1140 key includes the following feature improvements:

- Virtual backhitch (transaction write with sync)
- Single wrap backhitchless flush (large transaction writes with sync)
- Backhitchless backspacing (ANSI file writes)

The TS1140 function is improved with Recursive Accumulating Backhitchless Flush (RABF) and the addition of a new same wrap backhitchless flush (SWBF) function that extends virtual backhitch effectiveness for large files.

Read-ahead feature

On sequential reads, the tape drive automatically runs the read-ahead feature and fills the buffer with data sequentially beyond the target block (*N*). If one of these blocks is a target of the next command (such as *N*+200), it is already in the buffer. Therefore, transfer is fast. The drive then automatically fills the buffer sequentially with data past *N*+200.

The 3592-EH7 supports read ahead of approximately 1,000 MB of compressed data from tape. When the drive processes a command to locate or read a block, the drive automatically continues to stream down the tape and to read ahead until the data buffer is full, which allows subsequent `Locate` or `Read` commands to be fulfilled from the data buffer at faster speeds, rather than requiring access to the tape. The drive outperforms competitive drives, which stop and wait for the next command, with this unique functionality.

Data safe mode

The 3592 EH7 supports data safe mode. Data safe mode is controlled by the application and prevents inadvertent overwrite. Data safe mode treats the tape volume that is mounted as a WORM drive and prevents inadvertent overwrite. Data safe mode is set by the application or host system.

Drive mechanical and electrical reliability

The mechanism of the TS1140 tape drive is specified at a mean-cycles-between-failure rate of 300,000 cycles. The mechanism contains special mechanical and electrical features to prevent damage to the media on power-down or reset. These features also prevent the dropping of the leader pin or other thread failures during similar interruptions. It also tolerates high vibration and shock environments without data loss or degraded operation.

End-to-end logical block protection

The TS1140 supports end-to-end logical block protection, which allows the host application to append a cyclic redundancy check (CRC) to logical blocks. Data is validated in the following manner:

- Validated by the drive as data is received, and written to the media
- Validated by the drive and application on the data read-back
- Validated at full line speed with the **VERIFY** command

The write path is independent from the read path:

- Can read without a CRC and can write with a CRC
- Can write without a CRC and can read with a CRC
- Can write without a CRC and can read without a CRC
- Can write with a CRC and can read with a CRC

The CRC is generated and checked at the tape drive end and the application end. The CRC might be checked at any point along the path, at the host bus adapter (HBA), or at any system or software that is in the middle.

On write, the CRC is appended to each logical block at the source and the CRC is checked at any point that you want in the transfer. If corruption occurs, the exact location of the problem is known. The logical block with the CRC is validated before it is committed to media. Data is known to be on tape in uncorrupted form before the source data is released. The CRC is stored with each logical block.

On read, the logical block with the CRC is read and validated. The CRC is transferred with each logical block. The CRC is checked at any point that you want in the transfer, then the CRC is validated and removed at the destination. Logical block protection uses the **VERIFY** command as defined in the SCSI Stream Commands - 4 (SSC-4) standard, which adds new options to perform the following functions:

- Check that the protection information (that is, CRC) is the same protection information that is configured.
- Check the protection information (that is, CRC) and validate each logical block from the current position for the number (*n*) of files (that is, file marks (FMKS)):
  - No blocks are transferred over the SCSI interface.
  - Happens at the native data rate.
- Check that protection information (that is, CRC) validates on each logical block from the current position to EOD:
  - No blocks are transferred over the SCSI interface.
  - Happens at the native data rate.

2.5 IBM LTO Ultrium tape drives for the TS4500 tape library

Two LTO formats (Ultrium and Accelis) were introduced in 1997, and licenses for the technology were made available. Since then, the Accelis format was not actively pursued by manufacturers because it is apparent that the Ultrium format meets market needs. The three LTO sponsoring companies also took steps to protect the client's investment by providing a roadmap up to LTO generation 12, which illustrates native capacity (see Figure 2-10). The three LTO sponsoring companies also established an infrastructure to enable compatibility between products. At the time of this writing, five generations were available.

Important: IBM, Hewlett-Packard, and Quantum reserve the right to change the information in this migration path without notice.

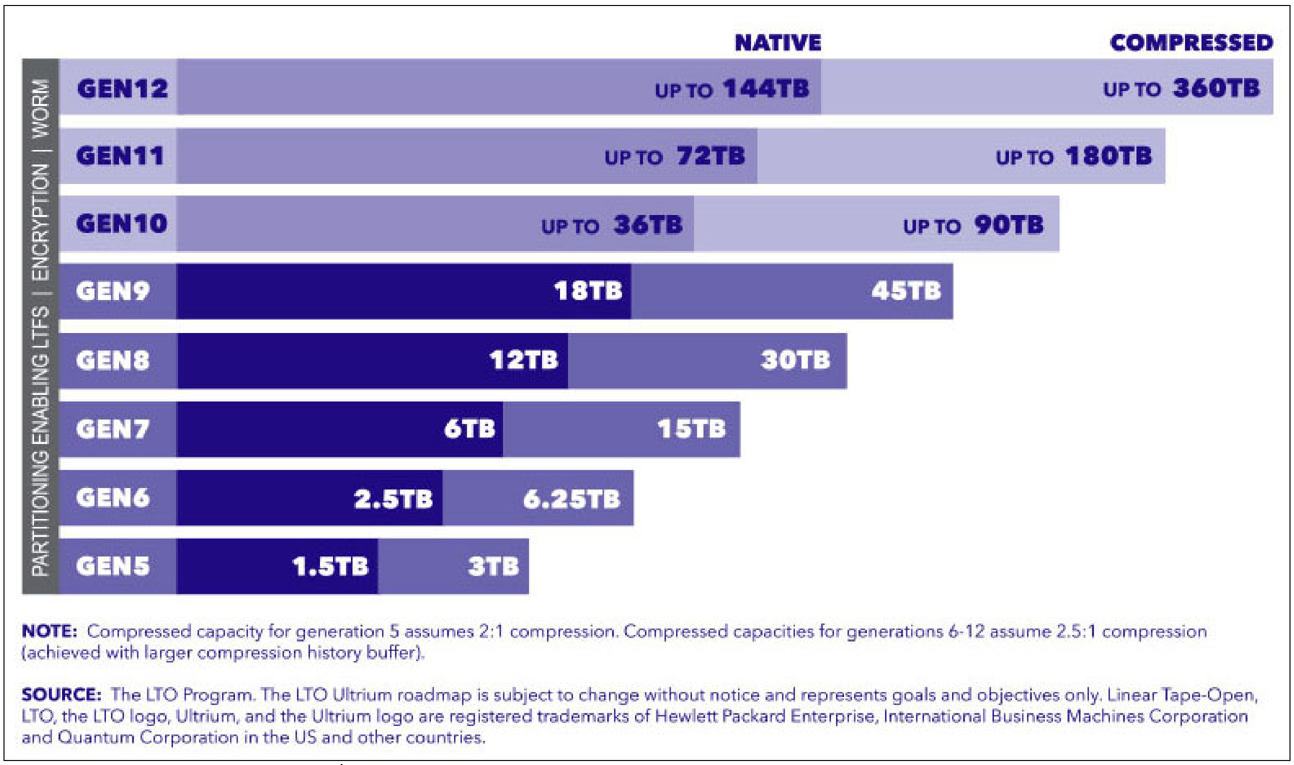

Figure 2-10 Ultrium Roadmap[1]

The LTO Ultrium compatibility investment protection is provided based on the following principles:

- Ultrium drive before LTO 7 are expected to read data from a cartridge in its own generation and at least the two previous generations
- Ultrium 9 and Ultrium 8 drive expected to read data from its own generation and from the immediate previous generation in the format of that generation
- An Ultrium drive is expected to write data to a cartridge in its own generation and to a cartridge from the immediate previous generation in the format of that generation

[1] For more information, see "What is LTO Technology?" at https://www.lto.org/what-is-lto/.

The three technology provider companies (IBM, HP, and Quantum) all made significant contributions of time and expertise to the definition of the LTO format specifications. All three companies have deep knowledge of clients' needs. They provided expert knowledge and engineering skill in the critical areas of magnetic recording technology, mechanism design, media materials, and cartridge design. This cooperative process created stronger LTO format definitions.

2.5.1 LTO core technology

Multichannel linear serpentine recording is at the core of the LTO formats. It enables an optimum balance of reliability and data integrity, performance, and high capacity. In the LTO recording format, data is written in tracks that run down the length of the tape.

The LTO Ultrium 9 format records data on 9860 tracks across the half-inch tape width. This linear recording format has a *serpentine characteristic*. The drive mechanism makes multiple passes from the beginning of the tape to the end of the tape and back to read or write the full capacity of the cartridge

In the LTO Ultrium 9 format, the 9860 tracks are split into four data bands of 2465 tracks each, and 32 read/write channels. The values for the Ultrium 8 and previous LTO generations are listed in Table 2-14.

Table 2-14 Data tracks, density, and channels

LTO generation	Data tracks	Tracks/band	Linear density	Read/write channels
Ultrium 1	384	96	124 Kbpi	8
Ultrium 2	512	128	188 Kbpi	8
Ultrium 3	704	176	250 Kbpi	16
Ultrium 4	896	224	328 Kbpi	16
Ultrium 5	1,280	320	368 Kbpi	16
Ultrium 6	2,176	544	385 Kbpi	16
Ultrium 7	3,584	896	485 Kbpi	32
Ultrium 8	6,656	1,664	524 Kbpi	32
Ultrium 9	9,860	2,465	545 Kbpi	32

Data is written to the innermost bands first to provide protection to the data that was recorded earliest in the process by writing it in the center, which is the most physically stable area on the tape. Data also is verified as it is written.

On the first pass of a round trip down the length of the tape and back, 32 tracks at LTO Ultrium 9 format are concurrently read or written (see the Read/write channels column in Table 2-14). At the end of the tape, the second pass of the round trip starts. The read/write heads are indexed and positioned over 32 new tracks, and the tape reverses direction back toward the beginning of the tape to complete the round trip. For the next round trip, the heads again are indexed to a new position over a new group of 32 tracks.

Because track densities (as shown in Table 2-14 on page 131) are high and because the tape is subject to lateral movement as it is moved, for performance and data integrity, the read/write heads must always be positioned precisely over the correct tracks. This positioning is accomplished by using the *timing-based servo* technique. This technique makes it possible to use high track densities (now and in the future) without changing the format of the media. This technique allows data with media imperfections to be read.

In the LTO system, electronic signals are generated through the real-time reading of servo data bands that are pre-recorded on the LTO tape. These signals enable the servo system to dynamically control the positioning of the read/write heads across the width of the tape. Similar magnetically based, track-following servo systems are used successfully in tens of thousands of tape drives that are in use today, such as the IBM TS1100 tape drive (IBM 3592).

The LTO formats also use advanced error correction codes for data integrity. These systems automatically correct most cross-track errors and provide data correction even if a full track is lost. Data is further protected by the demarcation of bad areas of the tape (for example, where servo signals are unreliable) and through dynamically rewriting bad blocks.

Cartridge memory is embedded in the LTO cartridges to record usage and error information. A noncontacting radio frequency (RFI) module, with nonvolatile memory capacity of 32,640 bytes for Ultrium 9, 16,320 bytes for Ultrium 8, 7 and 6, and 8192 bytes for Ultrium 5 and 4, provides storage and retrieval.

> **Note:** The TS4500 supports four LTO tape drives, the HD2-compatible models of the LTO-9, LTO-8, LTO-7, LTO-6, and LTO-5.

Interleaved recording

The LTO drive uses an interleaved, serpentine, longitudinal recording format that is similar to the method that is used in 3592 drives. The first set of 16 or 32 data tracks is written from near the physical beginning of the tape to near the physical end of the tape. The head then repositions to the next set of tracks for the return. This process continues until all tracks are written and the tape is full.

The format of the recording of the data and servo tracks is defined as part of the LTO specification to meet the requirement for interchange among implementations of different manufacturers.

Servo tracks

Servo tracks or *bands* enable the accurate positioning of the tape drive head over the data track, which ensures that the head does not stray onto an adjacent track. They are necessary to support high-data densities on the tape where the tracks are extremely close together. The servo bands are written when the cartridge is manufactured before the cartridge is usable for data storage and retrieval. If the servo bands are erased, the tape becomes unusable.

2.5.2 Data compression

The LTO Consortium created a superior data compression technique that is known as *LTO Data Compression* (LTO-DC). Although an excellent data compression algorithm, adaptive lossless data compression (ALDC) is not optimized for incompressible data, such as encrypted or previously compressed data.

For incompressible data, it is best not to apply any data compression algorithm, but rather to pass the input data directly to the compressed data stream (pass-through). ALDC might be preferable, and at other times, a simple pass-through is better, because of the variations in data. For example, if ALDC-based data compression is used, it is best if all segments of incompressible data are recorded without expansion by using a pass-through technique.

Figure 2-11 shows the LTO-DC data compression technique that uses the two schemes.

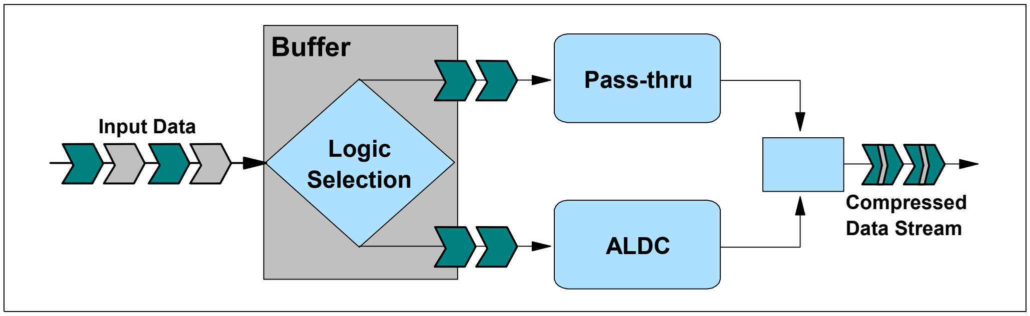

Figure 2-11 LTO-DC block diagram

The ability to swap schemes between ALDC and a pass-through mode gives a tape drive the power to automatically adapt to the incoming data stream.

No standardization of when to swap modes (scheme swap) when data is compressed was specified by LTO-DC. LTO-DC was approved by Ecma International as the Streaming Lossless Data Compression (SLDC) standard. For more information, see the Standard ECMA-321 Streaming Lossless Data Compression Algorithm - (SLDC) web page on the Ecma International website.

Compression technique: LTO uses the SLDC technique for compression. The IBM 3592 tape drive also uses the SLDC compression technique.

Because no standardization is specified, all vendor implementations might perform scheme swapping differently. What is specified and tested is that the resultant compressed data stream is decompressible by the defined set of LTO-DC rules. This capability enables interchange between drives from multiple vendors. Each vendor’s Ultrium drive can read and decompress the LTO-DC streams of the other vendors’ Ultrium drives.

2.5.3 Ultrium tape media

The IBM Ultrium cartridges are distinguished by color:

- Ultrium 6 cartridge is black
- Ultrium 7 cartridge is purple
- Ultrium 8 cartridge is burgundy
- Ultrium 9 cartridge is green

The IBM WORM cartridges are two-tone cartridges with a platinum bottom.

The top is like the normal LTO Ultrium cartridges. For the third generation of IBM WORM cartridges, the top is slate-blue. The Ultrium 6 WORM cartridge is black with a platinum bottom. The Ultrium 7 WORM cartridge is purple with a platinum bottom, The Ultrium 8 WORM is burgundy with platinum bottom and the Ultrium 9 WORM is green with platinum bottom.

Data cartridge

The Ultrium tape format specification is the implementation of LTO that is optimized for high capacity and performance with outstanding reliability, in a stand-alone or automated environment. The Ultrium cartridge uses a larger single-reel design (see Figure 2-12) and ½-inch tape to provide ultra-high storage capacity. The tape is extracted from the cartridge by the tape drive through a leader pin and is wound onto a take-up reel that is contained within the drive. This design is focused on client requirements for high capacity and performance and is ideally suited for backup, restore, and archive applications.

Ultrium drive technology is intended to meet the needs of the enterprise on a roadmap, or migration path, that extends well into the future. The Ultrium tape format established a new benchmark for large volume backup and archive options.

Figure 2-12 LTO-7 data cartridge (purple) and WORM cartridge (purple with platinum bottom)

In addition to standard read/write data cartridges, WORM cartridges are available. IBM WORM data cartridges are two tones to distinguish them from other data cartridges. Each WORM cartridge is the same color as the same generation of data cartridge on the top, but it is gray on the bottom. All generations of cartridges contain ½-inch, dual-coat, metal-particle tape. The three types of cartridges feature the following capacities:

- LTO-9 and LTO-9 WORM tape cartridges have a native data capacity of 18 TB (16.37 TiB) (45 TB [40.92 TiB] at 2.5:1 compression).
- LTO-8 and LTO-8 WORM tape cartridges have a native data capacity of 12 TB (10.91 TiB) (30 TB [27.93 TiB] at 2.5:1 compression).
- LTO-7 media formatted to M8 format, have a native data capacity of 9 TB (8.18 TiB) (22.5 TB [20.46 TiB] at 2.5:1 compression).
- LTO-7 and LTO-7 WORM tape cartridges have a native data capacity of 6 TB (5.45 TiB) (15 TB [13.32 TiB] at 2.5:1 compression).
- LTO-6 and LTO-6 WORM tape cartridges have a native data capacity of 2.5 TB (2.27 TiB) (6.25 TB [5.68 TiB] at 2.5:1 compression).
- LTO-5 and LTO-5 WORM tape cartridges have a native data capacity of 1.5 TB (1.36 TiB) (3 TB [2.73 TiB] at 2:1 compression).

The IBM WORM cartridges are two-tone cartridges with a platinum bottom. The top is like the normal LTO Ultrium cartridges.

WORM tape format

Beginning with LTO Ultrium format generation 3, WORM functionality provides nonerasable, nonrewritable operation with tape media. WORM is for long-term, tamper-resistant record retention. LTO Ultrium 9, 8, 7, 6, 5, and 4 drives provide the same WORM capability.

The format specification for WORM for LTO Ultrium generations includes low-level encoding in the cartridge memory, and the encoding is mastered into the servo pattern as part of the manufacturing process. This encoding prevents tampering.

Data can be appended at the end of a WORM cartridge to which data was previously written, which allows the full use of the high-capacity tape media.

Bar code labels

Each LTO data, cleaning, and diagnostic cartridge that is processed by the TS4500 tape library must bear a bar code label. The label contains the following information:

- A volume serial (VOLSER) number that you can read
- A bar code that the library can read

> **Note:** Client-printed bar code labels are not recommended for use on the TS4500. Certain bar code labels that were previously successfully scanned by the TS3500 are not readable by the TS4500. Issues with reading media labels are most often issues with those labels that were printed by the client.

When read by the library's bar code reader, the bar code identifies the cartridge's VOLSER to the tape library. The bar code also tells the library whether the cartridge is a data, cleaning, or diagnostic cartridge.

In addition, the bar code includes the two-character media-type identifier L*x*, where *x* equals 1, 2, 3, 4, 5, 6, 7, 8, 9, T, U, V, W, X. or Y identifies the cartridge as an LTO cartridge. M8 media type denoted an LTO 7 cartridge formatted for LTO 8 use only. The possible values for the second character are described in the following list:

- 1 indicates that the cartridge is the first generation of its type.
- 2, 3, 4, 5, 6, 7, M8, 8, or 9 indicates that the cartridge is the second, third, fourth, fifth, sixth, seventh, eighth, or ninth generation of its type
- T indicates that the cartridge is a generation 3 WORM cartridge
- U indicates that the cartridge is a generation 4 WORM cartridge
- V indicates that the cartridge is a generation 5 WORM cartridge
- W indicates that the cartridge is a generation 6 WORM cartridge
- X indicates that the cartridge is a generation 7 WORM cartridge
- Y indicates that the cartridge is a generation 8 WORM cartridge
- Z indicates that the cartridge is a generation 9 WORM cartridge

Figure 2-13 shows a sample bar code label for the LTO-8 tape cartridge.

Figure 2-13 View of the LTO-8 tape cartridge label

Apply the following guidelines whenever you use LTO bar code labels:

- Use only bar code labels that are approved by IBM.
- Do not reuse a label or reapply a used label over an existing label.
- Before you apply a new label, remove the old label by slowly pulling it at a right angle to the cartridge case.
- Use peel-clean bar code labels that do not leave a residue after they are removed. If glue residue is on the cartridge, remove it by gently rubbing it with your finger. Do not use a sharp object, water, or a chemical to clean the label area.
- Examine the label before you apply it to the cartridge. Do not use the label if voids or smears are in the printed characters or bar code. An application's inventory operation takes much longer if the bar code label is not readable.
- Remove the label from the label sheet carefully. Do not stretch the label or cause the edges to curl.
- Position the label within the recessed bar code label area.
- With light finger pressure, smooth the label so that no wrinkles or bubbles exist on its surface.
- Verify that the label is smooth and parallel, and it has no roll-up or roll-over. The label must be flat to within 0.5 mm (0.02 in.) over the length of the label and have no folds, missing pieces, or smudges.
- Do not place other machine-readable labels on other surfaces of the cartridge. They might interfere with the ability of the bar code reader to read the bar code.

Bar code label specifications are listed on the IBM LTO Ultrium Cartridge Label Specification web page.

Cleaning cartridges

The IBM LTO Ultrium Cleaning Cartridge (which is known as the *universal cleaning cartridge*) and the LTO Ultrium Cleaning Cartridge are compatible with all LTO tape drives.

The cleaning cartridge also contains a cartridge memory device, which automatically tracks the number of times that it was used. Cleaning cartridges must be replaced after 50 cleaning cycles. The firmware in the drive detects whether a cleaning tape is expired.

The volume serial (VOLSER) number on the cleaning cartridge's bar code label must begin with `CLN`, or the library treats the cleaning cartridge as a data cartridge during an inventory.

> **Note:** The average number of mounts that a cleaning cartridge is mounted before it is marked for replacement is 50. Because the count for cleaning is based on the tape length that is used during the cleaning, this number can be more or fewer than 50 mounts.

Cartridge memory

Within the cartridge is the LTO Cartridge Memory (LTO-CM), which is a passive, contactless silicon storage device that is physically a part of the cartridge. The memory chip is also known as medium auxiliary memory (MAM). For more information, see the *IBM TotalStorage LTO Ultrium Tape Drive SCSI Reference (LTO-5 through LTO-9)*, GA32-0928-04

Information about the cartridge and tape is written to the LTO-CM. The LTO-CM is only accessible and used by the drive itself and contains no client data. The LTO-CM is serial Electronically Erasable Programmable Read-Only Memory (EEPROM) with read-only and rewritable areas. It is housed inside the cartridge casing at the left rear (label side) corner.

The LTO-CM is used to hold usage and error information about the cartridge, the media inside that cartridge, and the data on the media. The storage capacity of the Generation 9 LTO-CM is 32,640 bytes. The Generation 8, 7, and 6 LTO-CM is 16,320 bytes, double the 8,160-byte capacity of the Generation 5 and 4 LTO-CM. The Generation 1, 2, and 3 cartridges use a 4,096-byte LTO-CM. The LTO-CM is in the left rear corner of the cartridge. A copy of this information also is kept in the first data set within the user data area, and it is given the data set number zero.

Communication between the drive and the LTO-CM is performed by a low-level radio frequency (RF) field that is generated (in the IBM implementation) by the drive. The LTO-CM is nonvolatile storage that is updated by using the RF field. It requires no other power source. This type of technology has an expected shelf life of more than 30 years.

Although transparent to the user, keeping this type of information enhances the efficiency of the drive and the cartridge. Data and block locations are stored in memory. For example, the end-of-data location is stored so that when the tape is next loaded, the drive can use the fast locate function to move directly to the recording area and begin recording.

Storing data about the age of the cartridge, the number of times that it was loaded, and the number of errors that it accumulated helps determine the reliability of the cartridge. This data is of particular value if it is stored with the cartridge itself, so that whenever the cartridge is mounted on any host system, the history is accessible.

This product is not the first tape product where information was kept on the cartridge. However, previously it was written on the tape medium in a portion of the tape (that users were not allowed to access) before the beginning-of-tape marker, such as in the IBM 3590 tape drive.

Cartridges compatibility

The cartridges that are compatible with the LTO tape drives that are supported by the tape library are listed in Table 2-15.

Table 2-15 Cartridges that are compatible with LTO tape drives

Cartridge type	Drive type[a]				
	LTO 9	LTO 8	LTO 7	LTO 6	LTO 5
18 TB LTO 9 Data Cartridge (xxxxxxL9)	R/W	Not compatible	Not compatible	Not compatible	Not compatible
18 TB LTO 9 WORM Cartridge (xxxxxxLZ)	R/W	Not compatible	Not compatible	Not compatible	Not compatible
12 TB LTO 8 Data Cartridge (xxxxxxL8)	R/W	R/W	Not compatible	Not compatible	Not compatible
12 TB LTO 8 WORM Cartridge (xxxxxxLY)	R/W	R/W	Not compatible	Not compatible	Not compatible
9 TB LTO M8 Data Cartridge (xxxxxxxM8)[b]	Not compatible	R/W	Not compatible	Not compatible	Not compatible
6 TB LTO 7 Data Cartridge (xxxxxxL7)	Not compatible	R/W	R/W	Not compatible	Not compatible
6 TB LTO 7 WORM Cartridge (xxxxxxLX)	Not compatible	R/W	R/W	Not compatible	Not compatible
2.5 TB LTO 6 Data Cartridge (xxxxxxL6)	Not compatible	Not compatible	R/W	R/W	Not compatible
2.5 TB LTO 6 WORM Cartridge (xxxxxxLW)	Not compatible	Not compatible	R/W	R/W	Not compatible
1.5 TB LTO 5 Data Cartridge (xxxxxxL5)	Not compatible	Not compatible	Read only	R/W	R/W
1.5 TB LTO 5 WORM Cartridge (xxxxxxLV)	Not compatible	Not compatible	Read only	R/W	R/W
800 GB LTO 4 Data Cartridge (xxxxxxL4)	Not compatible	Not compatible	Not compatible	Read only	R/W
800 GB LTO 4 WORM Cartridge (xxxxxxLU)	Not compatible	Not compatible	Not compatible	Read only	R/W
400 GB LTO 3 Data Cartridge (xxxxxxL3)	Not compatible	Not compatible	Not compatible	Not compatible	Read only
400 GB LTO 3 WORM Cartridge (xxxxxxLT)	Not compatible	Not compatible	Not compatible	Not compatible	Read only

a. The drive rejects any command to load unsupported media and returns a sense key of 3 and an extra sense code/additional sense code qualifier of 30/01.

b. Uninitialized M8 media and preinitialized M8 media are supported by a minimum LTO 8 tape drive firmware level of HB82 and a minimum library firmware level 1.4.1.2. The use of M8 media with prior firmware levels might result in a permanent reduction in the capacity of the M8 media 9 - 6 TB (native).

2.5.4 Physical attachment

The IBM LTO tape drive family offers high capacity, performance, and technology for the midrange open systems environment. These tape drives are dual-ported tape drives that facilitate 8 Gbps Fibre Channel connectivity. LTO drives are differentiated by their machine type and model numbers. You can identify the LTO tape drives by the label at the rear of the drive's canister.

IBM LTO tape drives can connect to many types of servers. For more information about the server interface cards or host bus adapters (HBAs) that are supported by the LTO tape drives, see the IBM System Storage Interoperation Center (SSIC) web page.

The TS4500 LTO Ultrium drives feature a front bezel with a chevron fiducial. The buttons and display are the same as for all models of LTO tape drives. The front of the LTO-9 (3588 F9C) tape drive is shown in Figure 2-14.

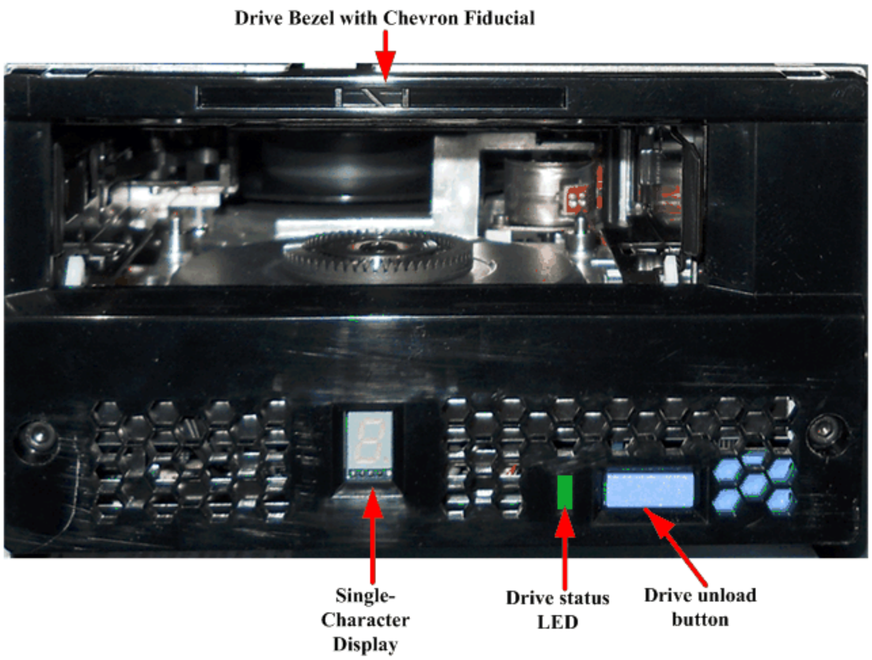

Figure 2-14 LTO Ultrium 9 (3588 F9C) front panel

The TS4500 LTO Ultrium back panel is the same back panel that is used in all LTO models (the exception is the model label). The back panel of the LTO-9 (3588 F9C) tape drive is shown in Figure 2-15.

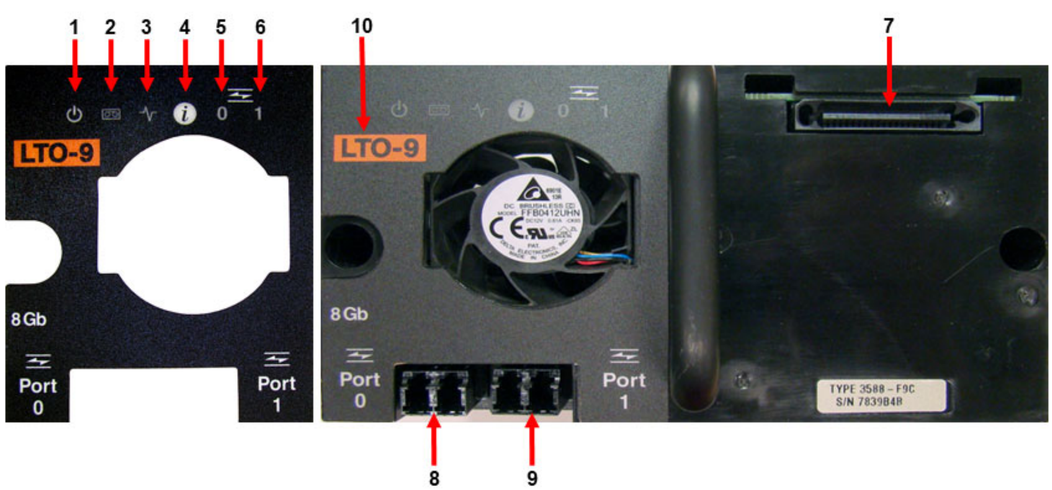

Figure 2-15 LTO Ultrium 9 Model F9C rear panel

The components on the rear panel and the normal status of the LEDs are listed in Table 2-16. The numbers in Table 2-16 refer to the numbers that are shown in Figure 2-15 on page 139.

Table 2-16 LTO Ultrium rear panel

Number	Description	LED color or description
1	Power status	Green
2	Drive status	Yellow or green
3	Library connection	Green
4	Information	Blue. Possible canister card issue. Drive fan starts running at high speed and turns on blue LED
5	Port 0 - Fibre Channel or SAS activity	Green or yellow
6	Port 1 - Fibre Channel or SAS activity	Green or yellow
7	TS4500 interface connector	TS4500 communication and power
8	Port 0 8 Gbps fiber connection or 12 Gbps SAS connection	Supports N_ports and NL_ports, which can autoconfigure between using the FC-AL protocol and the Direct Fabric attach protocol
9	Port 1 8 Gbps fiber connection or 12 Gbps SAS connection	Supports N_ports and NL_ports, which can autoconfigure between using the FC-AL protocol and the Direct Fabric attach protocol
10	LTO drive model	LTO-9, LTO-8, LTO-7, LTO-6, or LTO-5

2.5.5 Performance specifications

The performance specifications of the TS4500 LTO-9, LTO-8, LTO-7, LTO-6, and LTO-5 tape drives are listed in Table 2-17.

Table 2-17 Performance specifications for LTO tape drives

Performance parameter	Tape drive					
	LTO9 (SAS)	**LTO 9 (FC)**	**LTO 8 (FC)**	**LTO 7 (FC)**	**LTO 6 (FC)**	**LTO 5 (FC)**
Sustained data rate (native)[a]	400 MBps (L9 media)	400 MBps (L9 media)	360 MBps (L8 media)	L8 media not supported	L8 media not supported	L8 media not supported
	360 MBps (L8 media)	360 MBps (L8 media)	300 MBps (M8 media)	M8 media not supported	M8 media not supported	M8 media not supported
	M8 media not supported	M8 media not supported	300 MBps (L7 media)	300 MBps (L7 media)	L7 media not supported	L7 media not supported
	L7 media not supported	L7 media not supported	L6 media not supported	160 MBps (L6 media)	160 MBps (L6 media)	L6 media not supported
	L6 media not supported	L6 media not supported	L5 media not supported	140 MBps (L5 media)	140 MBps (L5 media)	140 MBps (L5 media)
	L5 media not supported	L5 media not supported	L4 media not supported	L4 media not supported	120 MBps (L4 media)	120 MBps (L4 media)
	L4 media not supported	L4 media not supported	L3 media not supported	L3 media not supported	L3 media not supported	80 MBps (L3 media)

Performance parameter	Tape drive					
	LTO9 (SAS)	LTO 9 (FC)	LTO 8 (FC)	LTO 7 (FC)	LTO 6 (FC)	LTO 5 (FC)
Sustained data rate (L6, L7, L8, and L9 media compressed at 2.5:1 compression; L5 and earlier media compressed at 2:1 compression)	900 MBps (L9 media)	750 MBps (L9 media)	750 MBps (L8 media)	L8 media not supported	L8 media not supported	L8 media not supported
	750 MBps (L8 media)	750 MBps (L8 media)	750 MBps (M8 media)	M8 media not supported	M8 media not supported	M8 media not supported
	M8 media not supported	M8 media not supported	750 MBps (L7 media)	750 MBps (L7 media)	L7 media not supported	L7 media not supported
	L7 media not supported	L7 media not supported	L6 media not supported	400 MBps (L6 media)	400 MBps (L6 media)	L6 media not supported
	L6 media not supported	L6 media not supported	L5 media not supported	280 MBps (L5 media)	280 MBps (L5 media)	280 MBps (L5 media)
	L5 media not supported	L5 media not supported	L4 media not supported	L4 media not supported	240 MBps (L4 media)	240 MBps (L4 media)
	L4 media not supported	L4 media not supported	L3 media not supported	L3 media not supported	L3 media not supported	160 MBps (L3 media)
Maximum sustained data rate (at maximum compression)	900 MBps (L9 media) 500 MBps (L8 media)	700 MBps (L9 media) 700 MBps L8 media)	750 MBps	750 MBps	745 MBps	745 MBps
Burst data rate for Fibre Channel drives	1200 MBps	800 MBps	800 MBps	800 MBps	800 MBps	800 MBps
Time to load, thread, and initialize a cartridge	17 seconds	17 seconds	15 seconds	15 seconds	12 seconds	12 seconds
Time to unload cartridge from load point	24 seconds	24 seconds	24 seconds	20 seconds	17 seconds	17 seconds
Average space record time from load point	45 seconds (L9 media)	45 seconds (L9 media	59 seconds (L8/M8/L7 media)	56 seconds (L7 media)	62 seconds (L6 media)	60 seconds (L5 media)

a. All sustained data rates depend on the entire data path (from data source to host system to tape drive).

2.6 IBM LTO Ultrium 9 tape drive (Model 3588 F9C, F9S, and S9C)

Ultrium 9 is the latest LTO generation, which was released in September 2021. The IBM LTO Ultrium 9 offerings represent significant improvements in capacity, performance, and reliability over the previous generation, Ultrium 8, while they still protect the client's investment in the previous technology.

The Ultrium 9 tape drive is a high-performance, high-capacity tape drive. The drive records data by using a linear serpentine recording format on half-inch tape that is housed within a cartridge. The data tracks are located by using preformatted servo tracks.

The Ultrium 9 tape drive provides the following improvements over the older Ultrium 8 models:

- The Ultrium 9 tape drive increases the native transfer speed to 400 MBps compared with 360 MBps for the Ultrium 8 tape drive.
- The Ultrium 9 FC tape drive compressed speed of 700 MBps for Fibre Channel interface and Ultrium 9 SAS tape drive compressed speed to 900 MBps for SAS interface.
- The Ultrium 9 tape drive dramatically increase the potential capacity of a cartridge in Ultrium 8 format.

3588 F9C, F9S and S9C tape drive features

The LTO-9 (3588 F9C, F9S and S9C) tape drive includes the following features:

- Native data transfer rate of up to 400 MBps
- Compressed data transfer rate of up to 700 MBps for model F9C/F9S and 900 MBps for model S9C
- LTO Ultrium 9 data and WORM tape cartridge native physical capacity of up to 18 TB
- 8 Gbps multi-mode Fibre Channel attachment option for 3588-F8C
- 8 Gbps single-mode Fibre Channel attachment option for 3588-F8S
- 12 Gbps dual-mode SAS attachment option for 3588-S9C
- 1024 MB internal data buffer
- Application-managed encryption support
- Library-managed encryption support
- 32 KB cartridge memory with Ultrium 9media
- Introducing Open Recommended Access Order (oRAO) to LTO 9 technology
- IBM Spectrum Archive Single Drive Edition (SDE), Library Edition (LE), and Enterprise Edition (EE) support
- Provides partitioning support, which, in conjunction with IBM Spectrum Archive, provides users with file-level access to tape data

The 3588-F9S drive with dual single mode interfaces is designed for the heavy demands of backup and archive tape storage applications that require high bandwidth over long distances.

The LTO Ultrium 9 supports integration in the IBM TS4500 library by using a new compact drive sled (see Figure 2-16).

Figure 2-16 LTO Ultrium 9 (3588 F9C) tape drive

Compatibility

Ultrium 9 drives can read and write on Ultrium 8; it cannot read previously Ultrium generation cartridges. Ultrium 9 drives also can read and write on Ultrium 8 cartridges at 12 TB capacity.

The drive also writes to tapes that can be read by other licensed LTO Ultrium 9 drives. In addition to the use of the IBM LTO Ultrium 8 data cartridge (with up to 12 TB capacity), the drive offers read/write capability for certified LTO Ultrium 8 tape cartridges.

Note: The LTO-9 tape drive can read or write to LTO9 and LTO8 format cartridges only.

The native data transfer rate when a data cartridge of another generation is processed is listed in Table 2-18.

Table 2-18 Native data transfer rate with various media

	18 TB (Ultrium 9)	**12 TB (Ultrium 8)**
Supported methods of operating	Read/write	Read/write
Native data rate Fibre Channel (MBps)	400	360
Native Capacity	18 TB (16.37 TIB)	12 TB (10.91 TiB)
2.5:1 compression capacity	45 TB (40.92 TiB)	30 TB (27.93 TiB)

The LTO Ultrium 9 is supported in the IBM TS4500 library by minimum library firmware version 1.7.0.3-A00 or higher.

3588 Field MES support

The following drive MES conversions are supported:

- 3588 F9A drive → 3588 F9C model upgrade
- 3588 F8A drive → 3588 F8C model upgrade
- 3588 F7A drive → 3588 F7C model upgrade
- 3588 F6A drive → 3588 F6C model upgrade
- 3588 F5A drive → 3588 F5C model upgrade

Connectivity

The LTO-9 tape drive family offers high capacity, performance, and technology for the midrange open systems environment. LTO-9 model F9C and F9S tape drives are dual-ported tape drives that facilitate 8 Gbps Fibre Channel connectivity. LTO-9 model S9S tape drive is a dual-ported tape drive that facilitated 12 Gbps SAS connectivity. These drives are differentiated from other LTO drives by their machine type and model numbers. You can identify the LTO tape drives by the logo at the front of the drive or by the label at the rear of the drive’s canister.

IBM LTO-9 tape drives can connect to many types of servers. For more information about server interface cards or host bus adapters (HBAs) that are supported by the LTO tape drives, see the IBM System Storage Interoperation Center (SSIC) web page.

Performance

The LTO Ultrium 8 tape drive uses 8960 data tracks to read and write to Ultrium 9 tape. These tracks are grouped in four servo bands. The high-bandwidth servo system features a low-mass servo to help more effectively track servo bands and improve data throughput with damaged media in less-than-optimal shock and vibration environments.

The native data transfer rate for Ultrium 9 SAS and FC tape drive model is 400 MBps. Compressed data rates can reach a peak rate of 700 MBps on the FC interface and 900 MBps on the SAS interface. IBM suggests the use of the IBM LTO Ultrium 9 18 TB data cartridge, which provides up to 45 TB of storage with a 2.5:1 compression ratio.

Open Recommended Access Order

A new technology y is introduced to IBM LTO 9 tape drive technology, Open Recommended Access Order (oRAO). This technology reduces tape data access times in LTO Ultrium Gen 9 tape technologies by up to 73%6.

oRAO enables tape control applications to accelerate the retrieval of a specific number of “files” from a single tape, which reduces the seek time between those “files” (see Figure 2-17).

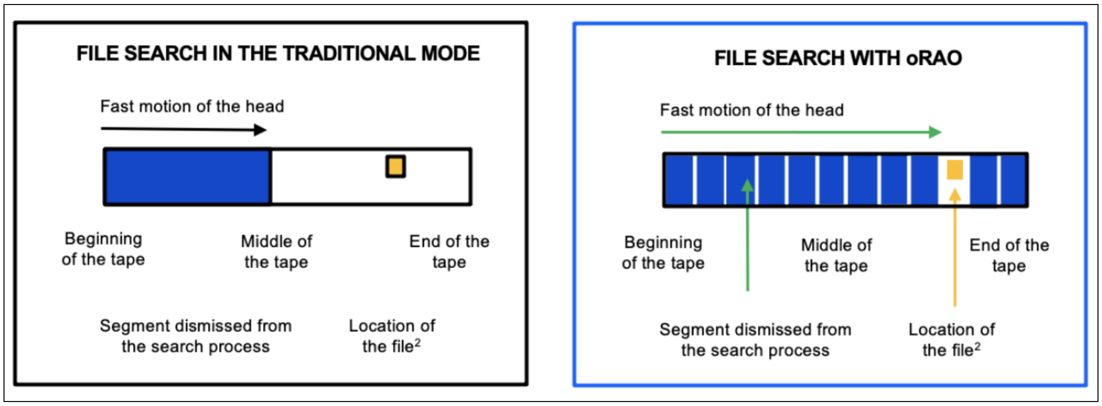

Figure 2-17 IBM LTO 9 Open Recommended Access Order

It was developed from IBM file access acceleration technology. It can add cyber resilience by optimizing the access times to recovery data. oRAO is a native drive function that supports compressed and uncompressed data and is available for LTO 9 generation technologies only.

Achieve up to 73% faster data access retrieval by using the IBM LTO 9 tape drive oRAO technology when working with larger volumes of data (see Figure 2-17).

Note: IBM LTO 9 oRAO is available on LTO-9 Full High tape drives only.

Dynamic speed matching

The LTO Ultrium 9 tape drive performs *dynamic speed matching* at one of 12 speeds to adjust the native data rate of the drive as closely as possible to the net host data rate (after it factors out data compressibility). This approach helps to reduce the number of backhitch repositions and improves throughput performance.

Backhitching is the condition that occurs when a data cartridge stops, reverses, and restarts motion. A backhitch is the result of a mismatch between the data rates of the connected server and the tape drive.

The data rates for the LTO Ultrium 9 drives are listed in Table 2-19.

Table 2-19 LTO-9 drive speed matching data rates

Performance parameters	Generation 9 media	Generation 8 media
Speed matching data rates in MBps	408 385 366 347 325 305 284 263 244 223 203 177	365.0 341.0 318.0 306.4 273.0 249.5 226.0 203.0 180.0 157.5 135.0 112.0

Encryption

The LTO Ultrium 9 tape drive family is encryption-capable and supports application-managed encryption (AME) at no charge on Fibre Channel or SAS tape drives. For library use, system-managed encryption (SME) and library-managed encryption (LME) are supported by the Transparent LTO Encryption feature code 1604 for TS4500. IBM Security Guardium Key Lifecycle Manager (GKLM) is required for encryption key management with LTO Ultrium 9 drives.

Statistical Analysis and Reporting System

The Ultrium 9 tape drives use Statistical Analysis and Reporting System (SARS) to help isolate failures between media and hardware. SARS uses the data cartridge performance history that is saved in the cartridge memory module and the drive performance history that is kept in the drive flash Electrically Erasable Programmable Read-Only Memory (EEPROM) to help determine the likely cause of the failure. SARS can cause the drive to request a cleaning tape, mark the media as degraded, and indicate that the hardware is degraded.

IBM Spectrum Archive (LTFS) software application

IBM Ultrium 9 tape drives are compatible with the IBM Spectrum Archive software application, and the underlying LTFS. LTFS uses the LTO media partitioning functionality, which is present on Ultrium 9, 8, 7, 6, and 5 cartridges.

LTFS provides a standard tape cartridge format at low cost that can be used without other database applications. LTFS presents tape media as though it were a disk file system. IBM Spectrum Archive supports the IBM LTO Ultrium 9, 8, 7, 6, and 5, and IBM TS1160, TS1155, TS1150, and TS1140 tape drives. IBM Spectrum Archive supports the LTO-7 initialized LTO-8 Type M cartridge. For more information, see “M8 format media” on page 148.

Tape as a storage medium offers many benefits. It is reliable, portable, low-cost, low-power, and high-capacity. However, tape is not easy to use. Tape has no standard format, and data often cannot be used without first copying the data to disk.

With LTFS, accessing data that is stored on an IBM tape cartridge became as easy and intuitive as the use of a USB flash drive. With LTFS file manager, reading data on a tape cartridge is as easy as dragging and dropping. Users can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape.

LTFS implements a true file system for tape. LTFS also supports library automation, including the ability to find data on a tape in a library without mounting and searching tape volumes.

IBM Spectrum Archive supports IBM tape automation in addition to the single drive edition. By using IBM Spectrum Archive Library Edition (LE), you can create a single file system mount point for a logical library that is managed by a single instance of the software, which runs on a single server. It also provides for caching of tape indexes and searching, querying, and displaying tapes' contents within an IBM tape library without the requirement to mount tape cartridges.

A TS4500 top-rack offering also can provide 10U of rack space on top of the library for Fibre Channel switches, tape data movers, or IBM Spectrum Archive nodes.

For more information about IBM Spectrum Archive and LTFS, see 2.11, "IBM Spectrum Archive" on page 162.

2.7 IBM LTO Ultrium 8 tape drive (Model 3588 F8C, F8S)

Ultrium 8 was released in November 2017. The IBM LTO Ultrium 8 offerings represent significant improvements in capacity, performance, and reliability over the previous generation, Ultrium 7, while they still protect the client's investment in the previous technology.

The Ultrium 8 tape drive is a high-performance, high-capacity tape drive. The drive records data by using a linear serpentine recording format on half-inch tape that is housed within a cartridge. The data tracks are located by using preformatted servo tracks.

The Ultrium 8 tape drive provides the following improvements over the older Ultrium 7 models:

- The Ultrium 8 tape drive increases the native transfer speed to 360 MBps compared with 300 MBps for the Ultrium 7 tape drive.
- The Ultrium 7 tape drive compressed speed of 700 MBps for Fibre Channel interface.
- The Ultrium 8 tape drive more than doubles the potential capacity of a cartridge in Ultrium 7 format.

3588 F8C, F8S tape drive features at a glance

The LTO-8 (3588 F8C, F8S) tape drive includes the following features:

- Native data transfer rate of up to 360 MBps
- Compressed data transfer rate of up to 700 MBps
- LTO Ultrium 8 data and WORM tape cartridge native physical capacity of up to 12 TB
- 8 Gbps multi-mode Fibre Channel attachment option for 3588-F8C
- 8 Gbps single-mode Fibre Channel attachment option for 3588-F8S
- 1024 MB internal data buffer
- Application-managed encryption support
- Library-managed encryption support
- Sixteen KB cartridge memory with Ultrium 8 media
- IBM Spectrum Archive Single Drive Edition, Library Edition, and Enterprise Edition support
- LTFS partitioning support

The 3588-F8S drive with dual single mode interfaces is designed for the heavy demands of backup and archive tape storage applications that require high bandwidth over long distances.

The LTO Ultrium 8 supports integration in the IBM TS4500 library by using a new compact drive sled.

The TS4500 LTO drive sled is shown in Figure 2-18.

Figure 2-18 LTO Ultrium 8 (3588 F8C)

Compatibility

Ultrium 8 drives can read and write on Ultrium 8 and Ultrium 7 cartridges only; it cannot read Ultrium 6 cartridges. Ultrium 8 drives also can read and write on Ultrium 7 cartridges at 9 TB capacity. For more information about M8 format, see “M8 format media” on page 148.

The drive also writes to tapes that can be read by other licensed LTO Ultrium 7 drives. In addition to using the IBM LTO Ultrium 7 data cartridge, with up to 6 TB capacity, the drive offers read/write capability for certified LTO Ultrium 7 tape cartridges.

> **Note:** The LTO-8 tape drive can read or write to LTO8 and LTO7 format cartridges only.

The native data transfer rate when a data cartridge of another generation is processed is listed in Table 2-20.

Table 2-20 Native data transfer rate with various media

	12 TB (Ultrium 8)	**9 TB M8 Format (Ultrium 7 media)**	**6 TB L7 Format (Ultrium 7)**
Supported methods of operating	Read/write	Read/write	Read/write
Native data rate Fibre Channel (MBps)	360	360	300
Native Capacity	12 TB (11.176 TiB)	9 TB (7.450 TiB)	6 TB (5.587 TiB)

	12 TB (Ultrium 8)	9 TB M8 Format (Ultrium 7 media)	6 TB L7 Format (Ultrium 7)
2.5:1 compression capacity	30 TB (28.61 TiB)	22.5 TB (20.954 TiB)	15 TB (13.64 TiB)

M8 format media

The LTO Program introduced a new capability with LTO-8 tape drives. The LTO-8 drives can write 9 TB (native) on a brand new LTO Ultrium 7 cartridge instead of 6 TB (native) as specified by the LTO 7 format. Such a cartridge is called an LTO 7 initialized LTO 8 Type M cartridge. These LTO 8 Type M cartridges are identifiable by using an automation bar code label ending with the last two characters *M8*.

Only new, unused LTO Ultrium 7 cartridges can be initialized as M8 cartridges. After a cartridge is initialized as M8, it cannot be changed back to L7. Initialized M8 cartridges can only be written and read in an LTO 8 tape drive. LTO7 tape drives cannot read initialized M8 cartridges.

M8 cartridges can be purchased as either pre-initialized M8 data cartridges or uninitialized M8 data cartridges. For either option, the bar code label is included. However, the uninitialized M8 data cartridge must first be initialized in tape libraries that support the automatic initialization of uninitialized M8 cartridges while under the control of ISV applications that recognize the M8 bar code label.

Note: M8 WORM cartridges are not supported.

A tape cartridge is initialized when it is first loaded into a compatible tape drive and data is written by the ISV application at the beginning of tape. The tape drive establishes the density of the media at that time.

Important: If an uninitialized M8 cartridge is not initialized in a tape library that supports uninitialized M8 cartridges, then the cartridge can be initialized at the L7 density even if the bar code label states M8. This situation can occur with the usage of non-TS4500 tape libraries, stand-alone LTO 7 tape drives, stand-alone LTO 8 tape drives, down-level LTO 8 tape drive firmware, down-level TS4500 tape library firmware, or down-level ISV software that does not recognize the M8 label.

Cartridges should only be mounted in LTO 8 tape drives. M8 cartridges that are inadvertently initialized at the L7 density can continue to be read and written in LTO 7 and LTO 8 tape drives, but they remain limited to the 6 TB native capacity.

The TS4500 tape library requires a minimum firmware version of 1.4.1.2. The drive needs a minimum LTO 8 tape drive firmware version of HB82 to support uninitialized M8 cartridges and to support for pre-initialized M8 cartridges.

3588 Field MES support

The following drive MES conversions are supported:

- 3588 F8A drive > 3588 F8C model upgrade
- 3588 F7A drive > 3588 F7C model upgrade
- 3588 F6A drive > 3588 F6C model upgrade
- 3588 F5A drive > 3588 F5C model upgrade
- 3592 E08 drive > 3592 EH8 model conversion

Connectivity

The LTO-8 tape drive family offers high capacity, performance, and technology for the midrange open systems environment. These tape drives are dual-ported tape drives that facilitate 8 Gbps Fibre Channel connectivity. These drives are differentiated from other LTO drives by their machine type and model numbers. You can identify the LTO tape drives by the logo at the front of the drive or by the label at the rear of the drive's canister.

IBM LTO-8 tape drives can connect to many types of servers. For more information about server interface cards or host bus adapters (HBAs) that are supported by the LTO tape drives, see the IBM System Storage Interoperation Center (SSIC) web page.

Performance

The LTO Ultrium 7 tape drive uses 3584 data tracks to read and write to Ultrium 8 tape. These tracks are grouped in four servo bands. The high-bandwidth servo system features a low-mass servo to help more effectively track servo bands and improve data throughput with damaged media in less-than-optimal shock and vibration environments.

The native data transfer rate for Ultrium 8 tape drives is 360 MBps. Compressed data rates can reach a peak rate of 800 MBps on the FC interface. IBM suggests the use of the IBM LTO Ultrium 8 12 TB data cartridge, which provides up to 30 TB of storage with a 2.5:1 compression ratio.

Dynamic speed matching

The LTO Ultrium 8 tape drive performs *dynamic speed matching* at one of 12 speeds to adjust the native data rate of the drive as closely as possible to the net host data rate (after it factors out data compressibility). This approach helps to reduce the number of backhitch repositions and improves throughput performance. *Backhitching* is the condition that occurs when a data cartridge stops, reverses, and restarts motion. A backhitch is the result of a mismatch between the data rates of the connected server and the tape drive.

The data rates for the LTO Ultrium 8 drives are listed in Table 2-21.

Table 2-21 LTO-8 drive speed matching data rates

Performance parameters	Generation 8 media	Generation 7 media M8 Format	Generation 7 media
Speed matching data rates in MBps	365.0 341.0 318.0 306.4 273.0 249.5 226.0 203.0 180.0 157.5 135.0 112.0	365.0 341.0 318.0 306.4 273.0 249.5 226.0 203.0 180.0 157.5 135.0 112.0	306.00 287.52 268.56 250.66 231.86 213.06 194.26 175.46 157.67 138.52 120.11 101.46

Encryption

The LTO Ultrium 8 tape drive family is encryption-capable and supports application-managed tape encryption at no charge on Fibre Channel tape drives. For library use, system-managed tape encryption and library-managed tape encryption are supported by the Transparent LTO Encryption features (FC 5901 for TS2900, FC 1604 for TS3500, and FC 5900 for all other libraries). IBM Security Guardium Key Lifecycle Manager is required for encryption key management with LTO Ultrium 8 drives.

Statistical Analysis and Reporting System

The Ultrium 8 tape drives use Statistical Analysis and Reporting System (SARS) to help isolate failures between media and hardware. SARS uses the data cartridge performance history that is saved in the cartridge memory module and the drive performance history that is kept in the drive flash Electrically Erasable Programmable Read-Only Memory (EEPROM) to help determine the likely cause of the failure. SARS can cause the drive to request a cleaning tape, mark the media as degraded, and indicate that the hardware is degraded.

IBM Spectrum Archive software application

IBM Ultrium 8 tape drives are compatible with the IBM Spectrum Archive software application, and the underlying LTFS. LTFS uses the LTO media partitioning functionality, which is present on Ultrium 8, 7, 6, and 5 cartridges.

LTFS provides a standard tape cartridge format at low cost that can be used without other database applications. LTFS presents tape media as though it were a disk file system. IBM Spectrum Archive supports the IBM LTO Ultrium 8, 7, 6, and 5, and IBM TS1155, TS1150, and TS1140 tape drives. IBM Spectrum Archive supports the LTO-7 initialized LTO-8 Type M cartridge. For more information, see “M8 format media” on page 148.

Tape as a storage medium offers many benefits. It is reliable, portable, low-cost, low-power, and high-capacity. However, tape is not easy to use. Tape has no standard format, and data often cannot be used without first copying the data to disk.

With LTFS, accessing data that is stored on an IBM tape cartridge became as easy and intuitive as using a USB flash drive. With LTFS file manager, reading data on a tape cartridge is as easy as dragging and dropping. Users can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape.

LTFS implements a true file system for tape. LTFS also supports library automation, including the ability to find data on a tape in a library without mounting and searching tape volumes.

IBM Spectrum Archive supports IBM tape automation in addition to the single drive edition. By using IBM Spectrum Archive Library Edition, you can create a single file system mount point for a logical library that is managed by a single instance of the software, which runs on a single server. In addition, it provides for caching of tape indexes and for searching, querying, and displaying tapes’ contents within an IBM tape library without the requirement to mount tape cartridges.

For more information about IBM Spectrum Archive and LTFS, see 2.11, “IBM Spectrum Archive” on page 162.

2.8 IBM LTO Ultrium 7 tape drive (Model 3588 F7C)

Ultrium 7 was released in late 2015. The IBM LTO Ultrium 7 offerings represent significant improvements in capacity, performance, and reliability over the previous generation, Ultrium 6, while they still protect the client's investment in the previous technology.

The Ultrium 7 tape drive is a high-performance, high-capacity tape drive. The drive records data by using a linear serpentine recording format on half-inch tape that is housed within a cartridge. The data tracks are located by using preformatted servo tracks.

The Ultrium 7 tape drive provides the following improvements over the older Ultrium 6 models:

- The Ultrium 7 tape drive increases the native transfer speed to 300 MBps compared with 160 MBps for the Ultrium 6 tape drive.
- The Ultrium 7 tape drive increases the compressed speed to 700 MBps for Fibre Channel compared with 400 MBps for the Ultrium 6 tape drive.
- The Ultrium 7 tape drive more than doubles the potential capacity of a cartridge in Ultrium 6 format.

3588 F7C tape drive features at a glance

The LTO-7 (3588 F7C) tape drive has the following characteristics:

- Native data transfer rate of up to 300 MBps
- Compressed data transfer rate of up to 750 MBps
- LTO Ultrium 7 data and WORM tape cartridge native physical capacity of up to 6 TB
- Eight Gbps Fibre Channel attachment option
- 1,024 MB internal data buffer
- Application-managed encryption support
- Library-managed encryption support
- Sixteen KB cartridge memory with Ultrium 7 media
- IBM Spectrum Archive Single Drive Edition, Library Edition, and Enterprise Edition support
- Linear Tape File System (LTFS) partitioning support

The LTO Ultrium 7 supports integration in the IBM TS4500 library by using a new compact drive sled.

The TS4500 LTO drive sled is shown in Figure 2-19.

Figure 2-19 LTO Ultrium 7 (3588 F7C)

Compatibility

In addition to reading and writing to LTO Ultrium 7 tape cartridges, the Ultrium 7 tape drives can read and write to LTO Ultrium 6 cartridges and read LTO Ultrium 5 cartridges. They cannot read Ultrium 4, Ultrium 3, Ultrium 2, or Ultrium 1 cartridges.

The drive also writes to tapes that can be read by other licensed LTO Ultrium 7 drives. In addition to using the IBM LTO Ultrium Data Cartridge with up to 6 TB capacity, the drive offers read/write capability for certified LTO Ultrium 7 tape cartridges.

Note: The LTO-7 tape drive cannot read or write LTO-1, LTO-2, LTO-3, or LTO-4 tape cartridges.

The native data transfer rate when a data cartridge of another generation is processed is listed in Table 2-22.

Table 2-22 Native data transfer rate with various media

	Generation 7 media	Generation 6 media	Generation 5 media
Supported methods of operating	Read/write	Read/write	Read only
Native data rate Fibre Channel (MBps)	300	160	140

3588 Field MES support

The following drive MES conversions are supported:

- 3588 F7A drive > 3588 F7C model upgrade
- 3588 F6A drive > 3588 F6C model upgrade
- 3588 F5A drive > 3588 F5C model upgrade

Connectivity

The LTO-7 tape drive family offers high capacity, performance, and technology for the midrange open systems environment. These tape drives are dual-ported tape drives that facilitate 8 Gbps Fibre Channel connectivity. These drives are differentiated from other LTO drives by their machine type and model numbers. You can identify the LTO tape drives by the logo at the front of the drive or by the label at the rear of the drive's canister.

IBM LTO-7 tape drives can connect to many types of servers. For a list of server interface cards or host bus adapters (HBAs) that are supported by the LTO tape drives, see the IBM System Storage Interoperation Center (SSIC) web page.

Performance

The LTO Ultrium 7 tape drive uses 3584 data tracks to read and write to Ultrium 7 tape. These tracks are grouped in four servo bands. The high-bandwidth servo system features a low-mass servo to help more effectively track servo bands and improve data throughput with damaged media in less-than-optimal shock and vibration environments.

The native data transfer rate for Ultrium 7 tape drives is 300 MBps. Compressed data rates can reach up to 750 MBps on the FC interface and 600 MBps on the serial-attached SCSI (SAS) interface. IBM suggests the use of the IBM LTO Ultrium 7 6 TB data cartridge, which provides up to 10 TB of storage with a 2.5:1 compression ratio.

Dynamic speed matching

The LTO Ultrium 7 tape drive performs *dynamic speed matching* at one of 12 speeds to adjust the native data rate of the drive as closely as possible to the net host data rate (after it factors out data compressibility). This approach helps to reduce the number of backhitch repositions and improves throughput performance. *Backhitching* is the condition that occurs when a data cartridge stops, reverses, and restarts motion. A backhitch is the result of a mismatch between the data rates of the connected server and the tape drive.

The data rates for the LTO Ultrium 7 drives are listed in Table 2-23.

Table 2-23 LTO-7 drive speed matching data rates

Performance parameters	Generation 7 media	Generation 6 media	Generation 5 media
Speed matching data rates in MBps	306.00 287.52 268.56 250.66 231.86 213.06 194.26 175.46 157.67 138.52 120.11 101.46	160.00 150.77 141.54 132.31 123.08 113.85 104.62 95.38 86.15 76.92 67.69 58.46 49.23 40.00	140.0 130.0 120.0 112.7 105.5 98.2 90.9 83.6 76.4 69.1 61.8 53.5 46.3 40.0

Encryption

The LTO Ultrium 7 tape drive family is encryption-capable and supports application-managed tape encryption at no charge on the SAS and Fibre Channel tape drives. In library use, system-managed tape encryption and library-managed tape encryption are supported by the Transparent LTO Encryption features (FC 5901 for TS2900, FC 1604 for TS3500, and FC 5900 for all other libraries). IBM Security Guardium Key Lifecycle Manager is required for encryption key management with LTO Ultrium 6 drives.

Statistical Analysis and Reporting System

The Ultrium 7 tape drives use Statistical Analysis and Reporting System (SARS) to help isolate failures between media and hardware. SARS uses the data cartridge performance history that is saved in the cartridge memory module and the drive performance history that is kept in the drive flash Electronically Erasable Programmable Read-Only Memory (EEPROM) to help determine the likely cause of the failure. SARS can cause the drive to request a cleaning tape to mark the media as degraded and to indicate that the hardware is degraded.

IBM Spectrum Archive software application

IBM Ultrium 7 tape drives are compatible with the IBM Spectrum Archive software application, and the underlying Linear Tape File System (LTFS). LTFS uses the LTO media partitioning functionality, which is present on Ultrium 7, 6, and 5 cartridges. LTFS provides a standard tape cartridge format at low cost that can be used without other database applications. LTFS presents tape media as though it were a disk file system. IBM Spectrum Archive supports the IBM LTO Ultrium 8, 7, 6, and 5, and IBM TS1155, TS1150, and TS1140 tape drives.

Tape as a storage medium offers many benefits. It is reliable, portable, low-cost, low-power, and high-capacity. However, tape is not easy to use. Tape has no standard format, and data often cannot be used without first copying the data to disk.

With LTFS, accessing data that is stored on an IBM tape cartridge became as easy and intuitive as using a USB flash drive. With LTFS file manager, reading data on a tape cartridge is as easy as dragging and dropping. Users can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape.

LTFS implements a true file system for tape. LTFS also supports library automation, including the ability to find data on a tape in a library without mounting and searching tape volumes.

IBM Spectrum Archive supports IBM tape automation in addition to the single drive edition. By using IBM Spectrum Archive Library Edition, you can create a single file system mount point for a logical library that is managed by a single instance of the software, which runs on a single server. In addition, it provides for caching of tape indexes and for searching, querying, and displaying tapes' contents within an IBM tape library without the requirement to mount tape cartridges.

For more information about IBM Spectrum Archive and LTFS, see 2.11, "IBM Spectrum Archive" on page 162.

2.9 IBM LTO Ultrium 6 tape drive (Model 3588 F6C)

The IBM LTO-6 tape drive was released in 2012. At the time, IBM LTO-6 tape drive offerings represented significant improvements in capacity and performance over the previous generation, LTO-5, while they protected the client's investment in the previous technology.

The 3588 F6C tape drive is a high-performance, high-capacity tape drive. The drive records data by using a linear serpentine recording format on half-inch tape that is housed within a cartridge. The data tracks are located by using preformatted servo tracks.

The 3588 F6C tape drive provides the following improvements over the older LTO models:

- The 3588 F6C tape drive increases the native transfer speed to 160 MBps compared with 140 MBps for the LTO-5 tape drive.
- The 3588 F6C tape drive increases the compressed speed to 400 MBps compared with 280 MBps for the LTO-5 tape drive.
- The 3588 F6C tape drive more than doubles the potential capacity of a cartridge in LTO-5 format.

3588 F6C tape drive features at a glance

The LTO-6 (3588 F6C) tape drive has the following characteristics:

- Native data transfer rate of up to 160 MBps
- Compressed data transfer rate of up to 400 MBps
- LTO-6 data and WORM tape cartridge native physical capacity of up to 2.5 TB
- Dual-ported 8 Gbps Fibre Channel attachment option
- 1,024 MB internal data buffer for full-high drives
- Application-managed encryption (AME) and Library-managed encryption (LME) support
- Sixteen KB cartridge memory with LTO-6 media
- IBM Spectrum Archive Single Drive Edition, Library Edition, and Enterprise Edition support
- IBM Linear Tape File System (LTFS) partitioning support

The LTO Ultrium 6 supports integration in the IBM TS4500 library by using a compact drive sled.

The TS4500 LTO drive sled is shown in Figure 2-20.

Figure 2-20 LTO Ultrium 6 (3588 F6C)

Compatibility

In addition to reading and writing to LTO-6 tape cartridges, the LTO 6 tape drives can read and write to LTO-5 tape cartridges and read LTO-4 tape cartridges.

Note: The LTO-6 tape drive cannot read or write to LTO-3, LTO-2, or LTO-1 tape cartridges.

The drive also writes to tapes that can be read by other licensed LTO-6 drives. In addition to using the IBM LTO Tape Cartridge with up to 2.5 TB capacity, the drive offers read/write capability for certified LTO-6 tape cartridges.

The native data transfer rate when a data cartridge of another generation is processed is listed in Table 2-24.

Table 2-24 Native data transfer rate with various media

	LTO-6 media	**LTO-5 media**	**LTO-4 media**
Supported methods of operating	Read/write	Read/write	Read only
Native data rate Fibre Channel (MBps)	160	140	120

3588 Field MES support

The following drive MES conversions are supported:

- 3588 F6A drive > 3588 F6C model upgrade
- 3588 F5A drive > 3588 F5C model upgrade

Connectivity

The LTO-6 tape drive family offers high capacity, performance, and technology for the midrange open systems environment. These tape drives are dual-ported tape drives that facilitate 8 Gbps Fibre Channel connectivity. These drives are differentiated from other LTO drives by their machine type and model numbers. You can identify the LTO tape drives by the logo at the front of the drive or by the label at the rear of the drive's canister.

IBM LTO-6 tape drives can connect to many types of servers. For a list of server interface cards or host bus adapters (HBAs) that are supported by the LTO tape drives, see the IBM System Storage Interoperation Center (SSIC) web page.

Performance

The LTO-6 tape drive uses 2,176 data tracks to read and write to an LTO-6 tape cartridge. These tracks are grouped in five servo bands. Like the LTO-5 tape drives, the high-bandwidth servo system features a low-mass servo to help more effectively track servo bands and improve data throughput with damaged media in less-than-optimal shock and vibration environments.

The native data transfer rate for LTO-6 tape drives is 160 MBps. Compressed data rates can reach 400 MBps. IBM suggests the use of the IBM LTO-6 2.5 TB tape cartridge, which provides up to 6.25 TB of storage with a 2.5:1 compression ratio.

Dynamic speed matching

The LTO-6 tape drive performs *dynamic speed matching* at one of 14 speeds to adjust the native data rate of the drive as closely as possible to the net host data rate (after it factors out data compressibility). This approach helps to reduce the number of backhitch repositions and improves throughput performance.

The data rates for the LTO-6 tape drives are listed in Table 2-25.

Table 2-25 LTO-6 tape drive speed matching data rates

Performance parameters	LTO-6 tape cartridge	LTO-5 tape cartridge	LTO-4 tape cartridge read only
Speed matching data rates in MBps	160.00 150.77 141.54 132.31 123.08 113.85 104.62 95.38 86.15 76.92 67.69 58.46 49.23 40.00	140.0 130.0 120.0 112.7 105.5 98.2 90.9 83.6 76.4 69.1 61.8 53.5 46.3 40.0	120.0 113.1 106.1 99.2 92.3 85.3 78.6 71.4 64.6 57.6 50.8 43.8 36.9 30.5

The LTO-6 tape drives are encryption capable, which means that they can convert data into a cipher that ensures data security. System-managed encryption and library-managed encryption and associated IBM Security Guardium Key Lifecycle Manager (GKLM) access are all available as a chargeable licensed key (FC 1604, Transparent LTO Encryption). A key is required to encrypt and decrypt the data. How a key is generated, maintained, controlled, and transmitted depends on the operating environment where the tape drive is installed.

Certain data management applications can manage keys. For an alternative solution, IBM provides a key manager that works with the keystore of your choice to perform all necessary key management tasks. No recovery exists for lost encryption keys.

Surface Control Guiding Mechanism

The *Surface Control Guiding Mechanism* is designed to guide the tape along the tape path in the LTO-6 tape drive. This method uses the surface of the tape rather than the edges to control tape motion.

LTO-6 tape drives use flangeless rollers so that the media can float naturally. The Surface Control Guiding Mechanism helps to reduce tape damage (especially to the edges of the tape) and tape debris that comes from the damaged edges and can accumulate in the head area.

Statistical Analysis and Reporting System (SARS)

The LTO-6 tape drives communicate with the TS4500 tape library through an internal Ethernet interface and use Statistical Analysis and Reporting System (SARS) to isolate failures between media and hardware. SARS uses the data cartridge performance history that is saved in the cartridge memory module and the drive performance history that is kept in the drive flash Electronically Erasable Programmable Read-Only Memory (EEPROM) to help determine the likely cause of the failure. SARS can cause the drive to request a cleaning tape to mark the media as degraded and to indicate that the hardware is degraded.

Media partitioning

Media partitioning is now available with the LTO-6 tape drive. Media partitioning allows for faster data access by splitting the cartridge into two media partitions. WORM media is not partitionable.

IBM Spectrum Archive software application

IBM Ultrium 6 tape drives are compatible with the IBM Spectrum Archive software application. IBM Spectrum Archive and the underlying Linear Tape File System (LTFS) use the LTO Ultrium 6 and Ultrium 5 media partitioning functionality. LTFS provides a standard tape cartridge format at low cost, and it can be used without other database applications. LTFS presents tape media as though it were a file system on a disk drive. LTFS supports only IBM IBM LTO Ultrium 8, 7, 6, and 5, and IBM TS1160, TS1155, TS1150, and TS1140 tape drives.

Tape as a storage medium has many benefits. Tape is reliable, portable, low-cost, low-power, and high-capacity. However, tape is not easy to use. It has no standard format, and data often cannot be used without first copying the data to disk.

With LTFS, accessing data that is stored on an IBM tape cartridge became as easy and intuitive as using a USB flash drive. With LTFS file manager, reading data on a tape cartridge is as easy as dragging and dropping. Users can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape.

LTFS implements a true file system for tape. LTFS also supports library automation, including the ability to find data on a tape in a library without mounting and searching tape volumes.

IBM Spectrum Archive supports IBM tape automation in addition to the single drive edition. By using IBM Spectrum Archive Library Edition, you can create a single file system mount point for a logical library that is managed by a single instance of the software, which runs on a single server. In addition, it provides for caching of tape indexes and for searching, querying, and displaying tapes' contents within an IBM tape library without the requirement to mount tape cartridges.

For more information about IBM Spectrum Archive and LTFS, see 2.11, “IBM Spectrum Archive” on page 162.

2.10 IBM LTO Ultrium 5 tape drive (Model 3588 F5C)

The IBM LTO-5 tape drive was released in 2010. At the time, IBM LTO-5 tape drive offerings represented significant improvements in capacity and performance over the previous generation, LTO-4, while they protected the client’s investment in the previous technology.

The LTO-5 tape drive is a high-performance, high-capacity tape drive. The drive records data by using a linear serpentine recording format on half-inch tape that is housed within a cartridge. The data tracks are located by using preformatted servo tracks.

LTO-5 tape drive features at a glance

The LTO-5 tape drive has the following characteristics:

- Native data transfer rate of up to 140 MBps
- Compressed data transfer rate of up to 280 MBps
- LTO-5 data and WORM tape cartridge native physical capacity of up to 1.5 TB
- Eight Gbps Fibre Channel attachment option
- 512 MB internal data buffer for full-high drives
- Application-managed encryption (AME) support
- Library-managed encryption (LME) support
- IBM Spectrum Archive Single Drive Edition, Library Edition, and Enterprise Edition support
- IBM Linear Tape File System (LTFS) partitioning support
- Eight KB cartridge memory with LTO-5 media

Connectivity

The LTO-5 tape drive family offers high capacity, performance, and technology for the midrange open systems environment. These tape drives are dual-ported tape drives that facilitate 8 Gbps Fibre Channel connectivity. These tape drives are differentiated from other LTO tape drives by their machine type and model numbers. You can identify the LTO tape drives by the logo at the front of the drive or by the label at the rear of the drive’s canister.

IBM LTO tape drives can connect to many types of servers. For a list of server interface cards or host bus adapters (HBAs) that are supported by the LTO tape drives, see the IBM System Storage Interoperation Center (SSIC) web page.

Compatibility

In addition to reading and writing to LTO-5 tape cartridges, the LTO-5 tape drives can read and write to LTO-4 tape cartridges and read LTO-3 tape cartridges. However, the LTO-5 tape drive cannot read LTO-2 or LTO-1 tape cartridges.

The drive also writes to tapes that can be read by other licensed LTO-5 tape drives. In addition to the use of the IBM LTO Tape Cartridge with up to 1.5 TB capacity, the drive offers read/write capability for certified LTO-5 tape cartridges.

The native data transfer rate when a tape cartridge of another generation is processed is listed in Table 2-26.

Table 2-26 Native data transfer rate with various media

	LTO-5 tape cartridge	LTO-4 tape cartridge	LTO-3 tape cartridge
Supported methods of operating	Read/write	Read/write	Read only
Native data rate Fibre Channel (MBps)	140	120	80

The LTO-5 tape drive family offers high capacity, performance, and technology for the midrange open systems environment. The LTO-6 tape drive (Model 3588 F6C) provides dual-ported 8 Gbps Fibre Channel connectivity for point-to-point or Fibre Channel Arbitrated Loop (FC-AL) attachment.

IBM LTO tape drives can connect to many types of servers. For a list of server interface cards or HBAs that are supported by the LTO tape drives, see the IBM System Storage Interoperation Center (SSIC) web page.

3588 Field MES support

The 3588 F5A drive > 3588 F5C model upgrade drive MES conversion is supported.

Performance

The LTO-5 tape drive uses 1,280 data tracks to read and write to the LTO-5 tape cartridge. These tracks are grouped in five servo bands. Like the LTO-4 tape drives, the high-bandwidth servo system features a low-mass servo to help more effectively track servo bands and improve data throughput with damaged media in less-than-optimal shock and vibration environments.

The native data transfer rate for LTO-5 tape drives is 140 MBps. Compressed data rates can reach 280 MBps. IBM advises the use of the IBM LTO-5 1.5 TB tape cartridge, which provides up to 3.0 TB of storage with a 2:1 compression ratio.

Dynamic speed matching

The LTO-5 tape drive performs *dynamic speed matching* at one of 14 speeds to adjust the native data rate of the drive as closely as possible to the net host data rate (after it factors out data compressibility). This approach helps to reduce the number of backhitch repositions and improves throughput performance.

The speed matching data rates for the LTO-5 tape drives are listed in Table 2-27.

Table 2-27 LTO-5 tape drive speed matching data rates

Performance parameters	LTO-5 tape cartridge	LTO-4 tape cartridge	LTO-3 tape cartridge
Speed matching data rates in MBps	140.0 130.0 120.0 112.7 105.5 98.2 90.9 83.6 76.4 69.1 61.8 53.5 46.3 40.0	120.0 113.1 106.1 99.2 92.3 85.3 78.6 71.4 64.6 57.6 50.8 43.8 36.9 30.5	80.0 76.1 72.3 68.4 64.6 60.8 56.9 53.1 49.2 45.4 41.5 37.7 33.8 30.0

Giant magneto-resistive head design

The LTO-5 tape drive uses the giant magneto-resistive (GMR) head with beveled contouring for reducing friction ("*stiction*"), especially with the smoother LTO-5 media type. This head design is well-proven in enterprise tape products, and it helps minimize contact, edge damage, debris accumulation, and wear on the tape as it moves over the read/write heads.

Surface Control Guiding Mechanism

The Surface Control Guiding Mechanism is designed to guide the tape along the tape path in the LTO-5 tape drive. This method uses the surface of the tape rather than the edges to control tape motion. LTO-5 tape drives use flangeless rollers so that the media can float naturally. The Surface Control Guiding Mechanism helps to reduce tape damage (especially to the edges of the tape) and tape debris, which comes from the damaged edges and can accumulate in the head area.

Statistical Analysis and Reporting System

The LTO-5 tape drive uses Statistical Analysis and Reporting System (SARS) to help isolate failures between media and hardware. SARS uses the data cartridge performance history that is saved in the cartridge memory module and the drive performance history that is kept in the drive flash Electronically Erasable Programmable Read-Only Memory (EEPROM) to help determine the likely cause of the failure. SARS can cause the drive to request a cleaning tape to mark the media as degraded and to indicate that the hardware is degraded. When a drive memory dump is taken from the drive, the IBM Support Center can determine whether the failure is in the tape drive or on the data cartridge.

Media partitioning

Media partitioning is now available with the LTO-5 drive. Media partitioning allows for faster data access by splitting the cartridge into two media partitions. WORM media is not partitionable.

IBM Spectrum Archive software application

The IBM Spectrum Archive software application and the underlying Linear Tape File System (LTFS) use the LTO Ultrium Generation 5 media partitioning functionality. LTFS provides a standard tape cartridge format at low cost, and it can be used without more database applications. LTFS presents tape media as though it were a file system on a disk drive. It supports only IBM LTO Ultrium 5 tape drives or higher-generation IBM LTO Ultrium tape drives.

Tape as a storage medium has many benefits. Tape is reliable, portable, low-cost, low-power, and high-capacity. However, tape is not easy to use. It has no standard format, and data often cannot be used without first copying the data to disk.

With LTFS, accessing data that is stored on an IBM tape cartridge became as intuitive as the use of a USB flash drive. With LTFS file manager, reading data on a tape cartridge is as easy as dragging and dropping. You can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape.

LTFS, which works on IBM LTO-8, 7, 6, 5, and IBM TS1160, TS1155, TS1150, and TS1140 tape drives, implements a true file system for tape. LTFS also supports library automation, including the ability to find data on a tape in a library without mounting and searching tape volumes.

IBM Spectrum Archive supports IBM tape automation in addition to the single drive edition. With IBM Spectrum Archive Library Edition, you can create a single file system mount point for a logical library that is managed by a single instance of the software, which is running on a single computer system. In addition, it provides for caching of tape indexes, and for searching, querying, and displaying tapes' contents within an IBM tape library without the requirement to mount tape cartridges.

For more information about IBM Spectrum Archive and LTFS, see 2.11, "IBM Spectrum Archive" on page 162.

2.11 IBM Spectrum Archive

IBM Spectrum Archive, a member of the IBM Spectrum Storage family, enables direct, intuitive, and graphical access to data that is stored in IBM tape drives and libraries by incorporating the Linear Tape File System (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). With IBM Spectrum Archive Enterprise Edition and IBM Spectrum Scale, a tape tier can now add savings as a low-cost storage tier. Network-attached unstructured data storage with native tape support by using LTFS delivers the best mix of performance and low-cost storage.

Key capabilities

IBM Spectrum Archive options can support small, medium, and enterprise businesses:

- Seamless virtualization of storage tiers
- Policy-based placement of data
- Single universal namespace for all file data
- Security and protection of assets
- Open, non-proprietary, and cross-platform interchange
- Integrated functionality with IBM Spectrum Scale

Benefits

IBM Spectrum Archive enables direct, intuitive, and graphical access to data that is 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 takes advantage of the low cost of tape storage and makes it as easy to use as drag-and-drop. Several IBM Spectrum Archive benefits are listed:

- Enable easy-as-disk access to single or multiple cartridges in a tape library
- Improve efficiency and reduce costs for long-term, tiered storage
- Optimize data placement for cost and performance
- Enable data file sharing without proprietary software
- Scale at a low cost
- Access and manage all data in stand-alone tape environments as though the data was on disk

2.11.1 Linear Tape File System (LTFS)

IBM addressed the growing storage needs of marketplace segments, such as media and entertainment, by introducing the first release of the Linear Tape File System (LTFS) in 2010. IBM developed IBM Spectrum Archive Single Drive Edition (SDE) to enable a self-describing cartridge that is based on dual partition in the LTO-5 technology. LTFS is a true file system that makes tape look and work like any removable media. IBM Spectrum Archive SDE is available as a no-charge download for single TS1160, TS1155, TS1150, TS1140, LTO-9, LTO-8, LTO-7, LTO-6, or LTO-5 tape drives from the IBM Fix Central web page.

LTFS is the first file system that works with LTO generation 9, 8, 7, 6, 5, TS1160, TS1155, TS1150, and TS1140 tape technology, to set a new standard for ease of use and portability for open systems tape storage. With this application, accessing data that is stored on an IBM tape cartridge is as easy and intuitive as using a USB flash drive. Tapes are self-describing, and you can quickly recall any file from a tape without reading the whole tape from beginning to end.

Furthermore, any LTFS-capable system can read a tape that is created by any other LTFS-capable system (regardless of the operating system platform). 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 displayed to you as a directory tree listing.
- More intuitive searches of cartridges and library content are now possible due to the addition of file tagging.

- Files can be moved to and from LTFS tape by using the familiar drag-and-drop technique that is common to many operating systems.
- Many applications that were written to use files on disk can now use files on tape without any modification.
- All standard File Open, Write, Read, Append, Delete, and Close functions are supported.

The annually archival data storage requirements continue to grow rapidly. The LTFS format is an ideal option for long-term archiving of large files that need to be easily shared with other individuals. This option is important because the LTO tape media that LTFS uses are designed for a 15-year to 30-year lifespan (depending on the number of read/write passes).

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

> **Important:** LTO Ultrium cartridges from earlier LTO generations (that is, LTO-1 through LTO-4) are not partitionable. Therefore, they cannot be used by LTFS. If LTO Ultrium 4 cartridges are used in an LTO Ultrium 5 drive to write data, the LTO-4 cartridge is treated like an unpartitioned LTO-5 cartridge. Even if an application can work with partitions, it is not possible to partition the LTO-4 media that is mounted at an LTO Ultrium 5 drive. Similarly, WORM cartridges of any generation cannot be used by LTFS because they cannot be partitioned.

The TS1160, TS1155, TS1150, and TS1140 tape drives are also supported by IBM Spectrum Archive SDE, IBM Spectrum Archive LE, and IBM Spectrum Archive EE.

Although LTFS presents the tape cartridge as a disk drive, the underlying hardware is still a tape cartridge and sequential in nature. Tape does not allow random access. Data is always appended to the tape, and no overwriting of files occurs. File deletions do not erase the data from tape but erase the pointers to the data.

So, although with LTFS, you can simultaneously copy two (or more) files to an LTFS tape, you get better performance if you copy files sequentially.

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

- Software in the form of an open source LTFS package
- Data structures that are created by LTFS on tape
- Hardware that consists of IBM LTO-9, LTO-8, LTO-7, LTO-6, LTO-5, TS1160, TS1155, TS1150, or TS1140 tape drives and tape media

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

LTFS partitioning support

Tape as a storage medium has many benefits:

- Reliable
- Portable
- Low cost
- Low power use
- High capacity

However, tape is not easy to use because it has no standard format, and data often cannot be used without first copying the data to a disk.

With LTFS, accessing data that is stored on an IBM tape cartridge became as easy and intuitive as the use of a USB flash drive. With LTFS file manager, reading data on a tape cartridge is as easy as dragging and dropping a file. Users can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape.

With IBM Spectrum Archive Single Drive Edition, you can create a single file system mount point for a logical library that is managed by a single instance of LTFS, which is running on a single computer system. In addition, it provides for caching of tape indexes, and for searching, querying, and displaying tapes' contents within an IBM tape library without the requirement to mount tape cartridges.

For more information about IBM Spectrum Archive SDE, see *IBM Linear Tape File System Installation and Configuration*, SG24-8090.

The TS1160, TS1155, TS1150, and TS1140 provide LTFS support with the following features:

- They provide the capability to configure up to four partitions.
- They support wrap-wise and longitudinal-wise partitioning.
- They support all non-WORM media formats.
- They provide `format` command support.
- Each partition can use a separate encryption method, or none.

The TS1160, TS1155, and TS1150 drives support partitioning in an identical manner to the TS1140, except the capacity of a specific partition scales up with the newly supported J5 and J5A format and improved performance:

- These drives support both the wrap-wise and longitudinal partitioning models and the same number of partitions.
- These drives support the partitioning of WORM media types to enable LTFS support on WORM media.
- These drives support a default wrap-wise partitioning model with minimal capacity loss at the expense of Accumulating Backhitchless Flush (ABF) capability within all partitions always.

Partitioning allows a volume to be split into multiple logical partitions, each of which can be read, navigated, written, erased, appended, updated, and managed as separate logical entities, with unique logical block sequences.

The primary user of this partitioning capability is LTFS, which partitions a volume into two logical partitions: an index partition and a data partition. The TS1160, TS1155, and TS1150 drives support both TS1140 style partitions.

The TS1160, TS1155, and TS1150 support more partitioning enhancements over the TS1140, specifically, both wrap-wise partitioning and longitudinal partitioning methods.

Restriction: The TS1160 does not support longitudinal partitioning.

The following characteristics apply in general to partitioning support on the TS1160 (3592 60E, 60F, and 60S), TS1155 (3592 55E and 55F), and TS1150 (3592 EH8):

- Partitioning is supported on media in the TS1160, TS1155, TS1150, and TS1140 logical formats.

  In JC media reuse, if you issue a `Format Medium` command, the system performs an implicit reformat to the drive format where the media is mounted; therefore, if TS1150 is used, it formats in TS1160 format if the media is in the TS1150 format.

  In JD media reuse, if you issue a `Format Medium` command, the system performs an implicit reformat to the TS1155 format, if the media is in the TS1150 format.

- Partitioning is only supported on unscaled R/W data and WORM media types that support writing in the TS1160, TS1155, TS1150, or TS1140 format.

  In scaled media, the `Format Medium` command is rejected.

  Attempts to scale partitioned media are accepted. As part of scaling, the volume is set to a single data partition cartridge.

  For more information about partitioning behavior, see IBM System Storage Tape Drive 3592 SCSI Reference, GA32-0068.

The following types of partitioning are available:

- Wrap-wise partitioning (which is used on LTO-9, LTO-8, 7, 6, and 5, TS1160, TS1155, TS1150, and TS1140).
- Longitudinal partitioning (maximum of two partitions) can be used on TS1155, TS1150, and TS1140 only.

Wrap-wise partitioning

LTO-9, 8, 7, 6, and 5 support two wrap-wise partitions. The TS1160, TS1155, TS1150, and TS1140 support four wrap-wise partitions.

Figure 2-21 shows wrap-wise partitioning. Consider the following points regarding wrap-wise partitioning:

- A maximum of four partitions are supported. Two or three partitions can be assigned, if preferred.
- A minimum of two wraps are allocated to a partition, regardless of the minimum selected capacity.
- The full length of the tape (LP3 to LP5) is always assigned to each partition.
- In general, two physical wraps between partitions are reserved as guard wraps. Therefore, a percentage of usable capacity might be lost, up to 3% for each partition boundary.
- RABF operations are performed in any partition if spare usable ABF wraps exist within a partition. In general, the last four wraps of a partition, or any partition that is smaller than four wraps, do not support RABF operations.

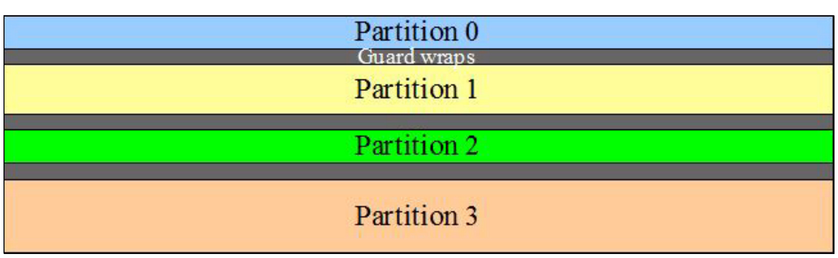

Figure 2-21 Wrap-wise partitioning

Longitudinal partitioning

The following conditions apply to longitudinal partitioning, which is supported on the TS1155, TS1150, and TS1140, TS1160:

> **Note:** For TS1160 that uses the new JE/JV/JM media types, wrap-wise partitioning is supported but longitudinal partitioning is not.
>
> For other media types (JD/JZ/JL and JC/JY/JK), the TS1160 drive continues to support both the wrap-wise and longitudinal partitioning models, similar to the E08 drive.

- A maximum of two partitions are supported.
- A minimum of 50 meters (164 ft.) are allocated to a longitudinal partition.
- The physical data wraps on the portion of tape that is assigned to the partition belong exclusively to each logical partition that is configured. Each partition starts from wrap 0 and ends on the last wrap.
- A guard gap between partitions is reserved to protect user data against systematic debris accumulation. The guardband is approximately 7 meters (22.9 ft.) and results in a capacity loss of less than 1%.
- RABF is performed within the boundaries of each partition, with the same wrap sequence as the base J5 RABF operation.
- Performance is slightly poorer due to less total ABF wrap length.
- As in non-partitioned media, the last four wraps' RABF cannot be used.
- Better performance is provided for random access because of shortened tape length for the partition.

Figure 2-22 shows longitudinal partitioning.

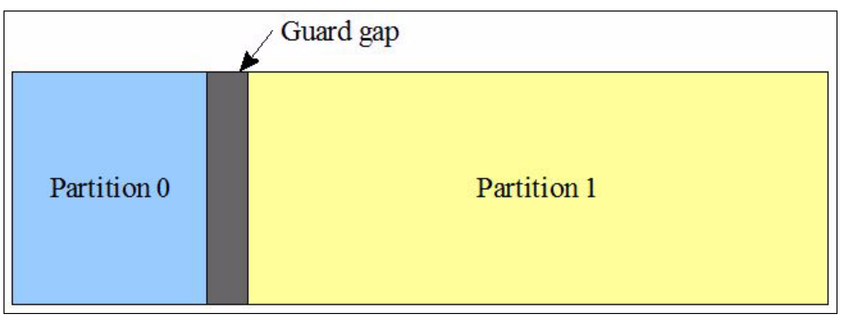

Figure 2-22 Longitudinal partitioning

2.11.2 IBM Spectrum Archive Editions

As shown in Figure 2-23, IBM Spectrum Archive is available in three editions that support small, medium, and enterprise businesses.

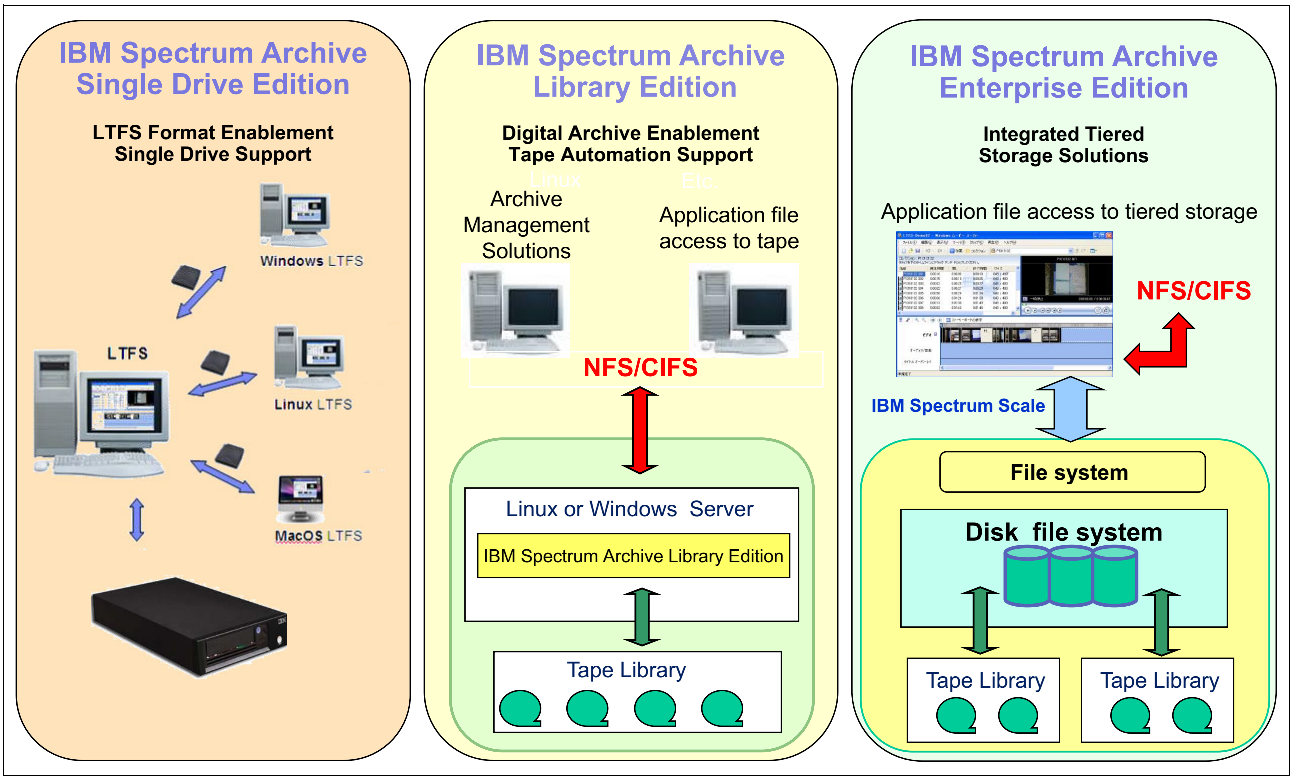

Figure 2-23 IBM Spectrum Archive SDE, LE, and EE

IBM Spectrum Archive Single Drive Edition

IBM Spectrum Archive Single Drive Edition (SDE) implements the LTFS format and allows tapes to be formatted as an LTFS volume. These LTFS volumes can then be mounted by using LTFS to allow users and applications direct access to files and directories that are stored on the tape. No integration with tape libraries exists in this edition. Access and manage all data in stand-alone tape environments as though the data was on disk.

IBM Spectrum Archive Single Drive Edition file system works with LTO generation 9, 8, 7, 6, 5, TS1160, TS1155, TS1150, and TS1140 tape technology to set a new standard for ease-of-use and portability for open systems tape storage. These tape drives support media partitioning, which allows their media to be divided in segments. Accessibility is improved because you can quickly locate data to retrieve or update.

LTFS can write files directly to tape media without any specific application. The tape drive shows up on the operating system as though it were a USB-attached drive. With LTFS file manager, reading data on a tape cartridge is as easy as dragging and dropping. Users can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape.

For more information about IBM Spectrum Archive SDE, see *IBM Linear Tape File System Installation and Configuration*, SG24-8090.

IBM Spectrum Archive Library Edition

IBM Spectrum Archive Library Edition (LE) extends the file manager capability of the IBM Spectrum Archive SDE. IBM Spectrum Archive LE is introduced with Version 2.0 of LTFS. Enable easy-as-disk access to single or multiple cartridges in a tape library.

LTFS is the first file system that works with IBM System Storage tape technology to optimize ease of use and portability for open systems tape storage. It manages the automation and provides operating system-level access to the contents of the library. IBM Spectrum Archive LE is based on the LTFS format specification, enabling tape library cartridges to be interchangeable with cartridges that are written with the open source SDE version of IBM Spectrum Archive. IBM Spectrum Archive LE supports most IBM tape libraries:

- TS2900 tape autoloader
- TS3100 tape library
- TS3200 tape library
- TS3310 tape library
- TS4300 tape library
- TS3500 tape library
- TS4500 tape library

IBM TS1160, TS1155, TS1150, and IBM TS1140 tape drives are supported on IBM TS4500, and TS1160, TS1150, and TS1140 are supported on the IBM TS3500 tape libraries only.

IBM Spectrum Archive LE enables the reading, writing, searching, and indexing of user data on tape and access to user metadata. *Metadata* is the descriptive information about user data that is stored on a cartridge. Metadata enables searching and accessing of files through the GUI of the operating system. IBM Spectrum Archive LE supports both Linux and Windows.

IBM Spectrum Archive LE provides the following product features:

- Direct access and management of data on tape libraries with LTO Ultrium 9, 8, 7, 6, 5 tape drives, and the TS1160, TS1155, TS1150, and TS1140 tape drives
- Tagging of files with any text, allowing more intuitive searches of cartridge and library content
- Exploitation of the partitioning of the media in the LTO-5 tape format standard
- One-to-one mapping of tape cartridges in tape libraries to file folders
- Capability to create a single file system mount point for a logical library that is managed by a single instance of LTFS and runs on a single computer system
- Capability to cache tape indexes and to search, query, and display tape content within an IBM tape library without needing to mount tape cartridges

The IBM Spectrum Archive LE offers the same basic capabilities as the SDE with more support of tape libraries. Each LTFS tape cartridge in the library appears as an individual folder within the file space. The user or application can navigate into each of these folders to access the files that are stored on each tape. The IBM Spectrum Archive LE software automatically controls the tape library robotics to load and unload the necessary LTFS volumes to provide access to the stored files.

Figure 2-24 shows how IBM Spectrum Archive LE presents the tapes in the library as folders.

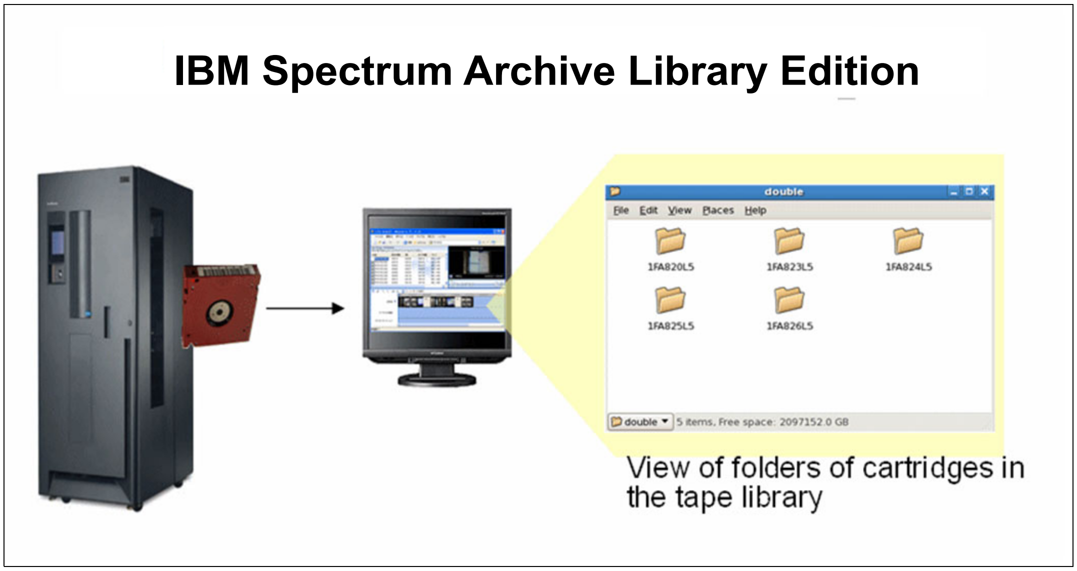

Figure 2-24 IBM Spectrum Archive LE view of tape folders

The following IBM tape libraries and tape autoloader support IBM Spectrum Archive LE:

- TS2900 tape autoloader
- TS3100 tape library
- TS3200 tape library
- TS3310 tape library
- TS4300 tape library
- TS3500 tape library
- TS4500 tape library

For more information about IBM Spectrum Archive LE, see *IBM Linear Tape File System Installation and Configuration*, SG24-8090.

IBM Spectrum Archive Enterprise Edition

IBM Spectrum Archive Enterprise Edition (EE) offers organizations an easy way to use cost-effective IBM tape drives and libraries within a tiered storage infrastructure. By using tape libraries instead of disks for tier 2 and tier 3 data storage (data that is stored for long-term retention), organizations can improve efficiency and reduce costs. In addition, IBM Spectrum Archive EE seamlessly integrates with the scalability, manageability, and performance of IBM Spectrum Scale, an IBM enterprise file management platform that enables organizations to move beyond adding storage to optimizing data management.

IBM Spectrum Archive Enterprise Edition offers the following benefits:

- Simplify tape storage with the IBM Linear Tape File System (LTFS) format in combination with the scalability, manageability, and performance of IBM Spectrum Scale
- Help reduce IT expenses by replacing tiered disk storage (tier 2 and tier 3) with IBM tape libraries
- Expand archive capacity by adding and provisioning media without affecting the availability of data that is already in the pool
- Add extensive capacity to IBM Spectrum Scale installations with lower media, less floor space, and lower power costs

IBM Spectrum Archive EE for the IBM TS4500, IBM TS3500, TS4300, and IBM TS3310 tape libraries provides seamless integration of IBM Spectrum Archive with IBM Spectrum Scale by creating an LTFS tape tier. You can run any application that is designed for disk files on tape by using IBM Spectrum Archive EE. 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. Improve efficiency and reduce costs for long-term, tiered storage.

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.

IBM Spectrum Archive EE supports IBM Linear Tape-Open (LTO) Ultrium 8, 7, 6, and 5 tape drives, and IBM System Storage TS1160, TS1155, TS1150, and TS1140 tape drives that are installed in TS4500 and TS3500 tape libraries or LTO Ultrium 9, 8, 7, 6, and 5 tape drives that are installed in the TS3310 tape library or LTO Ultrium 9, 8, 7, and 6 tape drives that are installed in the TS4300 tape library.

The use of IBM Spectrum Archive EE to replace disks with tape in tier 2 and tier 3 storage can improve data access over other storage solutions because it improves efficiency and streamlines the 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.

Figure 2-25 shows the hybrid storage solution with the integration of the IBM Spectrum Archive EE with IBM Spectrum Scale.

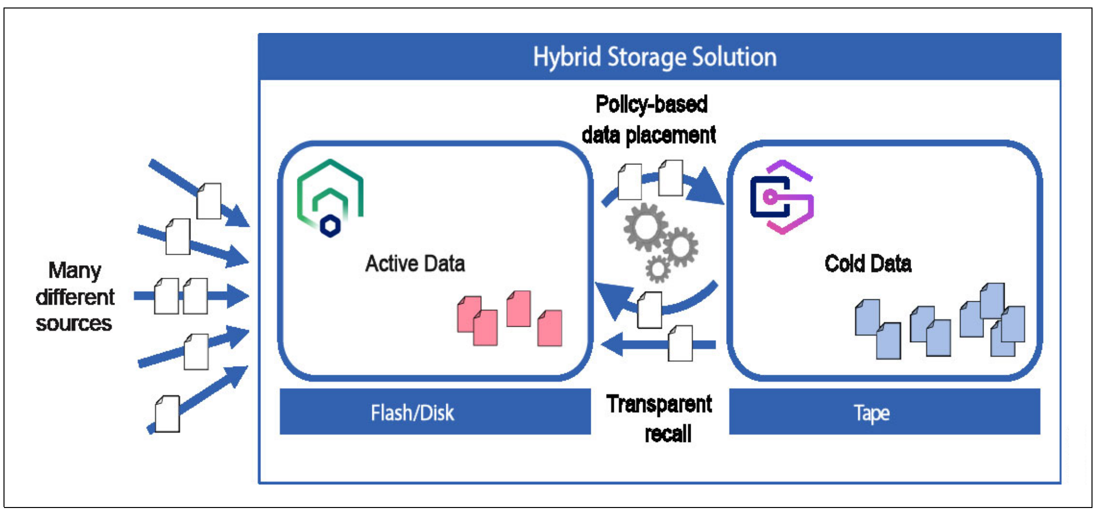

Figure 2-25 Integration of IBM Spectrum Scale and IBM Spectrum Archive Enterprise Edition

The seamless integration offers transparent file access in a continuous namespace. It provides the following capabilities:

- File-level write and read caching with a disk staging area
- Policy-based movement from disk to tape
- Creation of multiple data copies on different tapes
- Load balancing and high availability in multiple node clusters
- Data exchange on LTFS tape by using the import and export function
- Fast import of file namespace from LTFS tapes without reading data
- Built-in tape reclamation and reconciliation
- Simple administration and management

Note: For more information, see the following IBM Documentation web pages:

- IBM Spectrum Archive Enterprise Edition (EE)
- IBM Spectrum Archive Library Edition (LE)
- IBM Spectrum Archive Single Drive Edition (SDE)

3

Encryption

The encryption policies, methods, and software capabilities for the IBM TS4500 tape library are described in this chapter.

This chapter includes the following topics:

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

3.1 Tape encryption overview

The tape drives that are supported by the TS4500 tape library can encrypt data as it is written to a tape cartridge.

Encryption is performed at full line speed in the tape drive after compression. (Data is compressed more efficiently before it is encrypted.) This capability adds a strong measure of security to stored data without using processing power and without degrading performance.

3.1.1 Encryption-enabled tape drives

All of the tape drives that are supported by the TS4500 tape library are encryption-capable. *Encryption capability* means that they are functionally capable of performing hardware encryption, but this capability is not yet activated. To perform hardware encryption, the tape drives must be encryption-enabled. Encryption can be enabled on the tape drives through the TS4500 management graphical user interface (GUI).

Note: Transparent LTO Encryption FC, is required for library-managed encryption (LME) on Linear Tape-Option (LTO) tape drives. It is not required for application-managed encryption (AME). The current FC number can be found at IBM Documentation.

3.1.2 Encryption key management

Encryption involves the use of several kinds of keys in successive layers. How these keys are generated, maintained, controlled, and transmitted depends on the operating environment where the encrypting tape drive is installed. Specific data management applications, such as IBM Spectrum Protect (formerly called Tivoli® Storage Manager), can perform key management tasks.

For environments without such applications or environments where application-independent encryption is necessary, IBM provides a key manager to perform all necessary key management tasks. The suggested IBM Encryption Key Manager (EKM) for the TS4500 tape library and drives is IBM Security Guardium Key Lifecycle Manager (formerly called IBM Security Key Lifecycle Manager).

Note: Releases before Version 4.1.0 were IBM Security Key Lifecycle Manager; after Version 4.1.0, the name is IBM Security Guardium Key Lifecycle Manager.

IBM Security Guardium Key Lifecycle Manager is the IBM strategic platform for the storage and delivery of encryption keys to encrypt storage endpoint devices.

The IBM Security Guardium Key Lifecycle Manager can be used to provide encryption key management services for the encryption of data with encryption-capable drives. Host software has no direct knowledge of the key manager that is used.

IBM Security Guardium Key Lifecycle Manager serves data keys to the tape drive. You can use IBM Security Guardium Key Lifecycle Manager to create, back up, and manage the lifecycle of keys and certificates that an enterprise uses. You can manage encryption of symmetric keys, asymmetric key pairs, and certificates. IBM Security Guardium Key Lifecycle Manager provides a graphical user interface, command-line interface, and REST interface to manage keys and certificates.

For more information about the IBM Security Guardium Key Lifecycle Manager, see IBM Documentation.

3.2 Encryption policy

The *encryption policy* is the method that is used to implement encryption. It includes the rules that govern the volumes that are encrypted and the mechanism for key selection. How and where these rules are set up depends on the operating environment. For more information about each of the available methods, see 3.2.2, “Managing encryption on the TS4500” on page 176.

With the TS4500 tape library, the encryption policy is managed at the logical library level. The Logical Libraries page of the TS4500 management GUI is used to enable encryption for a logical library and modify the encryption method that is used. The Security page of the TS4500 management GUI is used to manage key servers and key labels.

> **Note:** In the tape storage environment, the encryption function on tape drives (desktop, standalone, and within libraries) is configured and managed by the client. It is not configured and managed by the IBM service support representative (SSR). In certain instances, SSRs are required to enable encryption at a hardware level when service access or service password-controlled access is required. Client setup support is from a field technical sales specialist (FTSS), client documentation, and software support for encryption software problems.

3.2.1 Encryption methods

The encryption methods for the TS1160, TS1155, TS1150, TS1140, LTO-9, LTO-8, LTO-7, LTO-6, LTO-5, and LTO-4 tape drives differ to some extent. The differences are described next. The following sections also include a brief description of encryption methods. In these sections, the term *Key Manager* (KM) is used to refer to IBM Security Guardium Key Lifecycle Manager and other key managers.

Symmetric key encryption

Encryption of data by using a symmetric key and algorithm is sometimes called *private key encryption* or *secret key*, which is not to be confused with the private key in an asymmetric key system. In a symmetric key system, the cipher key that is used for encrypting data is the same as the cipher key that is used for decryption.

The encryption and decryption ciphers can be related by a simple transformation on the key, or the encryption key and the decryption key can be identical. In the IBM Tape Encryption solution IBM Security Guardium Key Lifecycle Manager, the same encryption key is used for encryption and decryption of the data. This key is protected by an asymmetric key algorithm, and it is never available in clear text.

Symmetric key encryption is several orders of magnitude faster than asymmetric key encryption. Secret key algorithms can support encryption 1 bit at a time or by specified blocks of bits. The Advanced Encryption Standard (AES) supports 128-bit block sizes and key sizes of 128, 192, and 256. The IBM Tape Encryption solution uses the AES standard with a 256-bit key. Other well-known symmetric key examples are listed:

- Twofish
- Blowfish
- Serpent

- Cast5
- Data Encryption Standard (DES)
- Triple DES (TDES)
- International Data Encryption Algorithm (IDEA)
- Advanced Encryption Standard (AES)

Asymmetric key encryption

Another important method of encryption that is widely used today is referred to as *public/private key encryption* or *asymmetric encryption*. When this encryption methodology is used, ciphers are generated in pairs. The first key is used to encrypt the data. The second key is used to decrypt the data.

This technique was pioneered in the 1970s, and it represented a significant breakthrough in cryptography. The Rivest-Shamir-Adleman (RSA) algorithm is the most widely used public key technique. The power of this approach is a public key, which is used to encrypt the data.

This public key can be widely shared, and anyone who wants to send secure data to an organization can use its public key. The receiving organization then uses its private key to decrypt the data, which makes public/private key encryption useful for sharing information between organizations. This methodology is widely used on the internet today to secure transactions, including Secure Sockets Layer (SSL).

Asymmetric key encryption is much slower and more computationally intensive than symmetric key encryption. The advantage of asymmetric key encryption is the ability to share secret data without sharing the encryption key.

3.2.2 Managing encryption on the TS4500

A *key manager* is a software program that assists IBM encryption-enabled tape drives in generating, protecting, storing, and maintaining encryption keys. The encryption keys encrypt information that is being written to tape media (tape and cartridge formats), and they decrypt information that is being read from tape media.

The TS4500 tape library supports the IBM Security Guardium Key Lifecycle Manager. For more information, see the IBM Security Guardium Key Lifecycle Manager in IBM Documentation.

The key manager operates on a number of operating systems, including IBM z/OS, Linux, Sun Solaris, IBM AIX, and Microsoft Windows. It is a shared resource that is deployed in several locations within an enterprise. It can serve numerous IBM encrypting tape drives, or encrypting disk drives, regardless of where those drives are installed (for example, in tape library subsystems, which are connected to mainframe systems through various types of channel connections, or installed in other computing systems).

The key manager uses a *keystore* to hold the certificates and keys (or pointers to the certificates and keys) that are required for all encryption tasks. Refer to the appropriate documentation for detailed information about the key manager and the keystores that it supports.

The following methods are available to manage encryption in the TS4500 tape library:

- Application-managed encryption (AME)
- System-managed encryption (SME) for TS7700 z/OS
- Library-managed encryption (LME)

These methods differ in the following ways:

- Where the encryption policy engine resides
- Where key management occurs for your encryption solution
- How the key manager is connected to the drive

Your operating environment determines the best method for you.

Key management and the encryption policy engine can be in any of the environment layers that are shown in Figure 3-1.

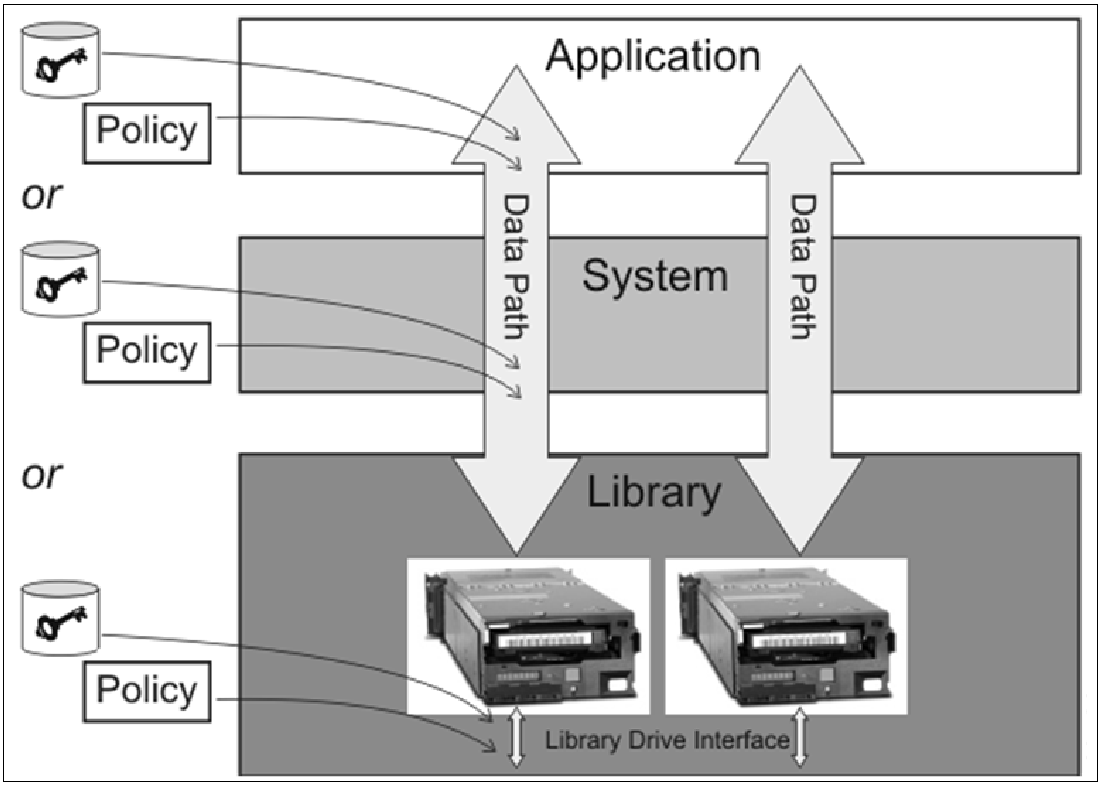

Figure 3-1 Possible locations for the encryption policy engine and key management

The application layer, for example, IBM Spectrum Protect, initiates the data transfer for tape storage.

The library layer is the TS4500 tape library, which contains an internal interface to each tape drive that is installed in the library.

3.2.3 Application-managed encryption

The application-managed encryption (AME) method is best in operating environments that run an application that already can generate and manage encryption policies and keys, such as IBM Spectrum Protect. Policies that specify when encryption is to be used are defined through the application interface. The policies and keys pass through the data path between the application layer and the encryption-capable tape drives.

Encryption is the result of interaction between the application and the encryption-enabled tape drive, and it is transparent to the system and library layers. Because the application manages the encryption keys, volumes that are written and encrypted with the application method can be read only by using the application-managed tape encryption method.

Note: The capability to use AME is not preset. The logical library must be set to use AME.

Application-managed tape encryption can use either of two encryption command sets:

- The IBM encryption command set that was developed for the key manager
- The T10 command set that was defined by the International Committee for Information Technology Standards (INCITS)

3.2.4 System-managed encryption

System-managed encryption (SME) is required for TS7700 support. Tape drives that attach to the TS7700 must be configured for system-managed encryption. The TS7700 can use the drives in this mode only, and it does not support library-managed or application-managed encryption.

After the TS7700 uses drives for encrypted physical tape volumes, it will place drives that are not correctly enabled for encryption offline to the subsystem.

System-managed encryption is best where the applications that write to or read from tapes are not capable of performing the key management that is required for application-managed encryption.

For IBM z Systems, encryption policies that specify when to use encryption can be set up in the z/OS Data Facility Storage Management Subsystem (DFSMS) or implicitly through each instance of an IBM device driver. Key generation and management are performed by an encryption key server. Policy controls and keys pass through the data path between the system layer and the encrypting tape drives. Encryption is transparent to the applications.

3.2.5 Library-managed encryption

Library-managed encryption (LME) is useful for encryption-enabled tape drives in an open-attached TS4500 tape library.

Note: The capability to use LME is not preset. The logical library must be set to use LME.

Key generation and management are performed by the key manager, which is a Java application that is running on a library-attached host. The keys pass through the library-to-drive interface. Therefore, encryption is transparent to the applications when it is used with certain applications, such as IBM Spectrum Protect.

Bar code encryption policies, which are set up through the TS4500 management GUI, can be used to specify when to use encryption. In such cases, policies are based on cartridge volume serial numbers (VOLSERs). Library-managed encryption also allows other options, such as the encryption of all volumes in a library, independently of bar codes. Key generation and management are performed by the key manager. Policy control and keys pass through the library-to-drive interface. Therefore, encryption is not apparent to the applications.

When it is used with certain applications, such as Symantec NetBackup or the EMC Legato NetWorker, library-managed encryption includes support for an internal label option. When the internal label option is configured, the encryption-enabled tape drive automatically derives the encryption policy and key information from the metadata that is written on the tape volume by the application.

Up to four library-managed encryption (LME) key paths per logical library are supported on the TS4500.

> **Note:** If you use LME and IBM device drivers that run on open systems platforms (AIX, Linux, Solaris, or Windows), information for bulk rekey is available in the *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130.
>
> When you use LME, an extra Ethernet cable must be attached, preferably to a different network switch. The extra cable is for redundancy and better backup job reliability.

The following components are required to use encryption:

- Encryption-enabled tape drive
- Keystore
- Key manager

3.2.6 Prerequisites for using encryption on the IBM TS4500 tape library

Certain hardware and software prerequisites must be met before you use encryption with the TS4500 tape library.

With the TS4500 tape library, encryption is managed at the logical library level. All encryption-enabled drives that are assigned to a logical library use the same method of encryption.

The rules for setting up encryption differ based on whether the library is installed with 3592 or LTO tape drives, and whether you use library-managed encryption (LME), application-managed encryption (AME), or system-managed encryption (SME). SME is not available for LTO drives, it is 3592 available only.

If the library contains 3592 tape drives, the following prerequisites must be met:

- IBM Security Guardium Key Lifecycle Manager is attached to the TS4500 and configured for LME.
- Tape drives are enabled for encryption from the Logical Libraries page of the TS4500 management GUI.
- If the tape drives connect to a TS7700, system-managed encryption is used.

If the library contains LTO tape drives, the following prerequisites must be met:

- Tape drives are enabled for encryption from the Logical Libraries page of the TS4500 management GUI.
- Transparent LTO Encryption is required for LTO tape drives if you use LME. For more information about the latest Feature Codes, see IBM Documentation.
- IBM Security Guardium Key Lifecycle Manager is the key manager when you use LME with LTO tape drives.

Add Encryption Key Server on the TS4500

Use the Encryption Key Servers page to manage which key servers use an encryption key.

To add an encryption key server or modify the server that is used, complete the following steps:

1. Click **Settings** → **Security** → **Encryption Key Servers**.
2. Select **Add Encryption Key Server**.
3. Enter the Server IP (if you use DNS name, ensure that you added the DNS servers in Network page) and the Port that is to be used.
4. Select the Security type: TLS1.2 or Proprietary. If TLS1.2 is selected (see Figure 3-2), a Key server certificate must be added to the TS4500 by using the folder browser. The file can contain a chain of certificates. The file also must be in base64 privacy enhanced mail (PEM) format. The maximum length is 255 characters.

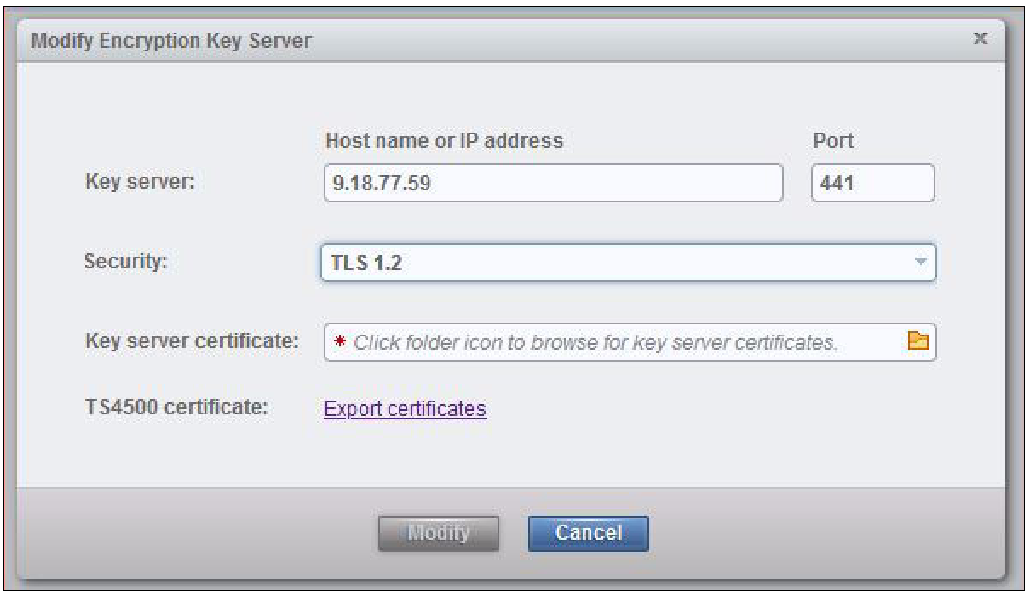

Figure 3-2 TLS1.2 option

5. Click **Modify**.

Encryption methods on the TS4500

Encryption is managed at the logical library level. All encryption-enabled drives that are assigned to a logical library use the same method of encryption. Enable encryption, or modify the method that is used, on the Logical Libraries page, as described in 4.1, “Integrated management console” on page 186.

To enable encryption or modify the method that is used, complete the following steps:

1. Select a logical library on the Logical Libraries page.
2. Select **Actions** → **Modify Encryption Method**.

3. Choose a method from the Encryption menu on the Modify Encryption Method window and click **Modify**, as shown in Figure 3-3.

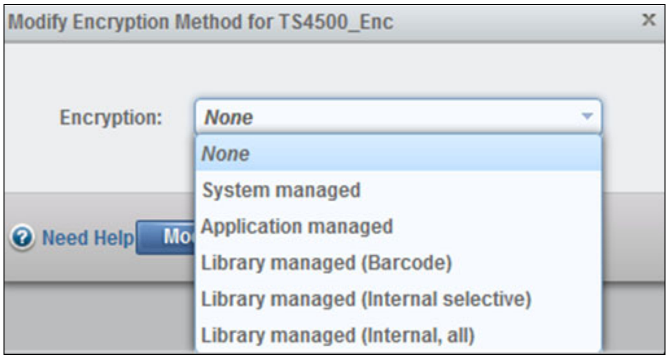

Figure 3-3 Modify Encryption Method window

For more information, see “Modify Encryption Method option” on page 293.

The following methods can be used for encryption:

- Application-managed encryption (AME)

  Use this method if the application generates and manages encryption policies and keys. Applications, such as IBM Spectrum Protect, can manage encryption.

- System-managed encryption (SME)

  Select this method of encryption if the library is attached to a TS7700 z/OS.

- Library-managed encryption (LME) by bar code

  Use this method to use the default key that is specified by the key manager for all VOLSER ranges. The encryption policy is specified based on cartridge volume serial numbers.

- Library-managed encryption (LME) by internal label selective encryption

  Use this method if you use Symantec NetBackup or the EMC Legato NetWorker. This encryption method encrypts cartridges with pool identifiers 1500 - 9999 (inclusive) by using keys that are specific to each pool.

  Labels for these keys are generated by the tape drive based on the pool identifier. For instance, key label `INTERNAL_LABEL_NBU_1505_A` is generated for a cartridge in pool `1505`. Go to **Settings** → **Security** → **Encryption Internal Label** and select the **Create mapping** tab to map these generated labels to the key-encrypting key labels that you want in the keystore of the Encryption Key Manager (EKM). All other cartridges remain unencrypted.

- Library-managed encryption (LME) by internal label all encryption

  Encrypt All Mode allows NetBackup to always request encryption and to specify the key labels to use. Certain ranges indicate that the default EKM key labels must be used and other ranges indicate that one or two key labels need to be constructed based on the pool ID.

  For NetWorker, Encrypt All Mode allows NetWorker to request encryption for all but two cases. The mode is the same as the Selective Encryption Mode, except where the Encryption Control Field (ECF) is invalid, out of range, or not provided. In this case, the drive generates a special “NOTAG” key label or labels. If the keystore has keys with this label, encryption occurs. However, the intended use of the “NOTAG” key label is to flag jobs that did not update their ECF for encryption.

If the “NOTAG” key is not in the keystore, the write fails and the job fails. This function allows the client to flag all jobs that were not altered for encryption.

Advanced Encryption Settings (For Service Use Only)

Advanced Encryption Settings allows only IBM Support personnel (under direction of the drive development team) to provide a work-around for an unforeseen problem or support a unique configuration. In some cases, this workaround can be done by using a PFE version of drive firmware without also creating a library firmware version.

This option is not intended for the customer to use without the guidance of IBM Support.

The settings are a full menu of potential operating modes for the drive that might override the behavior that is established by the method that is selected.Based on the use of these advanced encryption settings, attempting to match the method to the equivalent advanced setting is not needed. It is a library/drive firmware relationship that is not intended to be established.

Enabling Advanced Encryption Settings (Always encrypt) for a logical library on the TS4500

The following minimum code versions required:

- TS4500 Library Code: 1801-B00
- TS4500 CLI Tool: 1.8.0.1

If the logical library does not have encryption enabled, follow the steps tat are described in next. If the logical library has encryption enabled, see Step 3, on page 183.

Enable Library Managed Encryption for the Logical Library

Using TS4500 Web User Interface, complete the following steps from the Logical Libraries page:

1. Right-click the logical library and select **Modify Encryption**. The Modify Encryption Method menu opens (see Figure 3-4).

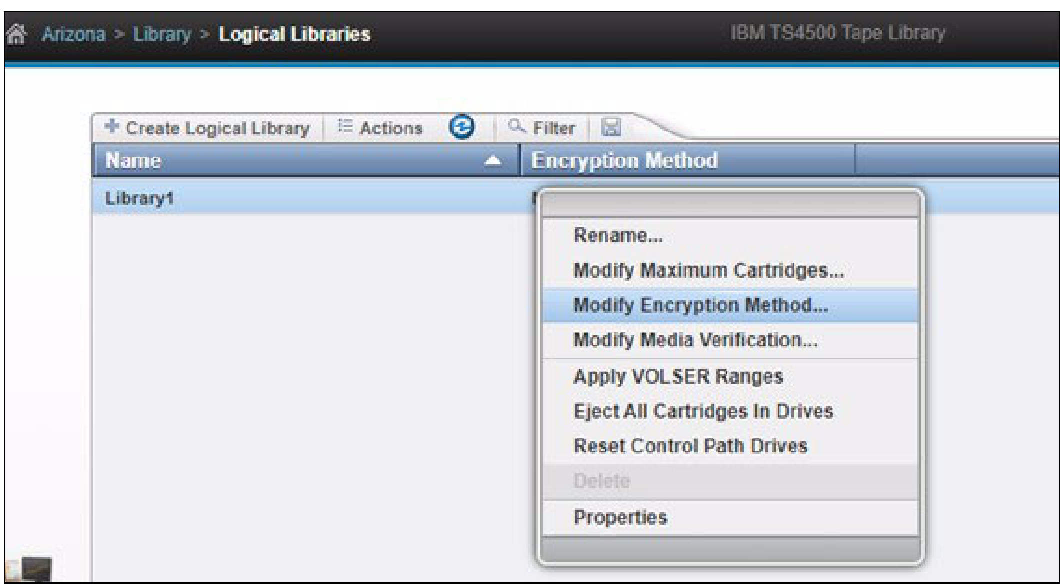

Figure 3-4 Logical Libraries Encryption Method options

2. In the Modify Encryption Method menu, select **Library Managed (Barcode)** and at least one key server to use with this logical library. Click **Modify** to save the settings, as shown in Figure 3-5.

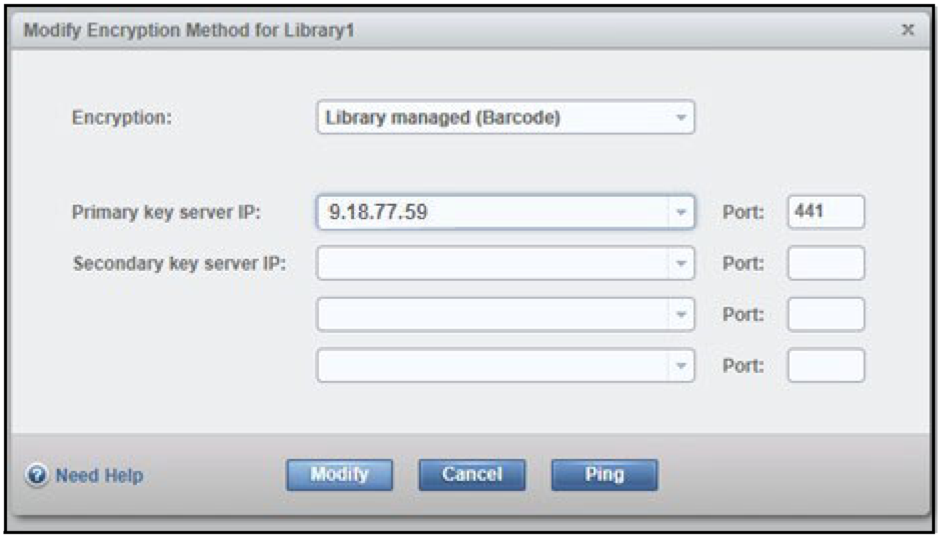

Figure 3-5 Modify Encryption Method for Library window

3. Set the advanced encryption (Always encrypt) for the logical library by using TS4500 CLI Tool v1.8.0.1:

   ```
   java -jar TS4500CLI.jar --ip address -u user -p  pwd
   --modifyAdvancedEncSettings Library1,TRUE,5,0,1
   ```

   Use the following advanced encryption settings:

   – Address is the IP address of the TS4500
   – User is a valid login for the TS4500 with Administrator role
   – User password for user is `pwd`
   – `Library1` is the name of the logical library
   – `TRUE` sets advanced encryption settings
   – `5` is the advanced-policy (always encrypt)
   – `0` is the density code (No advanced setting)
   – `1` is the keypath (No advanced setting)

4. Use the TS4500 CLI Tool and TS4500 Web User Interface to verify that the advanced encryption policy was configured.

5. Run the following CLI Tool command to verify advanced encryption settings for the logical library:

   ```
   java -jar TS4500CLI.jar --ip address_ip -u user -p pwd
   --viewAdvancedEncryptionSettings Library1
   ```

   An example of expected settings is shown in Figure 3-6.

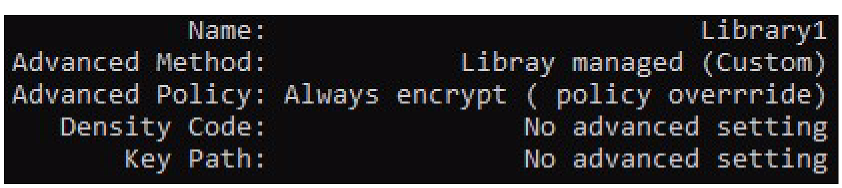

Figure 3-6 Example of expected settings

6. On the Logical Libraries page, verify that the encryption method for the logical library is `Library managed (Always encrypt)` (see Figure 3-7).

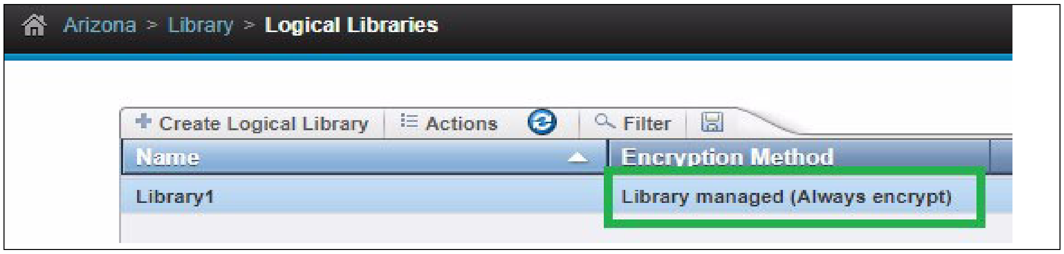

Figure 3-7 Verifying that the Encryption Method is Library managed (Always encrypt)

4

TS4500 management graphical user interface

The TS4500 provides advanced capabilities for integrated tape drive and media management, which are delivered within a “single-pane-of-glass” integrated management console (IMC). The TS4500 management graphical user interface (GUI) is based on a unified interface. This unified interface includes key features, which are used in other IBM storage solutions, to help guide storage administrators to complete critical tasks.

This chapter includes the following topics:

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

4.1 Integrated management console

The integrated management console (IMC) is a built-in platform for tools that can be used to locally manage the TS4500 tape library by using the TS4500 management GUI.

It is also used for local service activities and when enabled, remote support, and Call Home of the TS4500 tape library.

The IMC, which includes an LCD monitor and a keyboard with a touchpad, can be mounted on either end of the TS4500 tape library. A library controller card (LCC) and power source are required within the end frame or within the adjacent frame. Alternatively, you can use Feature Code (FC) 2737, IMC Separate Power Source, this moves the IMC PDU to an end frame and provides longer Ethernet cables to allow attachment of the IMC to an LCC anywhere in the library.

The TS4500 management GUI runs on a web browser in kiosk mode on the IMC. Kiosk mode means that the menu bar, address bar, and stop, and reload buttons of the browser are disabled. In addition, it is not possible to use bookmarks or multiple browser windows. Figure 4-1 shows the IMC.

Figure 4-1 The integrated management console

The IMC hardware components apart from the monitor and keyboard are located on the inside of the side panel, these being a Tiny PC, and two power supplies (one power supply for the monitor and the other for the Tiny PC). The Tiny PC models M93p, 9020, 7040, 7050, and 7060 are shown in Figure 4-2.

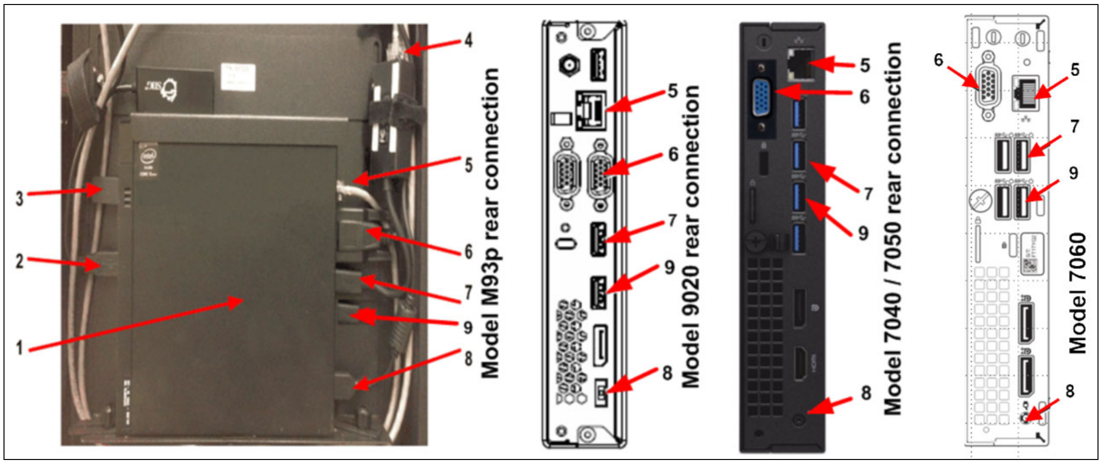

Figure 4-2 IMC Tiny PC

The components of the Tiny PC are listed in Table 4-1. The numbers correspond to the numbers that are shown in Figure 4-2.

Table 4-1 Tiny PC components

Number	Component
1	Tiny PC (models M93p, 9020, 7040, 7050, or 7060)
2	USB to Ethernet adapter
3	USB to Ethernet adapter
4	Ethernet cable to adapter
5	Ethernet cable to library controller card (LCC)
6	IMC video connector
7	USB to IMC keyboard and mouse
8	Power
9	DVD drive

4.2 Using the TS4500 management GUI

Before you connect to the TS4500 management GUI, the TS4500 tape library must first be installed and configured by an IBM Engineer. In addition, you must use one of the supported web browsers that are listed in Table 4-2 to access the TS4500 management GUI. To ensure that all of the functions of the management GUI are usable, enable cookies and JavaScript in your browser and disable the browser's function of blocking pop-up windows.

During installation, the IBM service support representative (SSR) configures the IMC by using network settings that are provided by the user. Static IP assignment is the default, but Dynamic Host Configuration Protocol (DHCP) can be selected. For static IP assignment, the user must provide the SSR with a listing of the values to enter in the network setup for the library IP address, subnet mask, and gateway.

4.2.1 Connecting to the management GUI

To connect to the management GUI, perform the following steps:

1. Open one of the supported web browsers, as listed in Table 4-2.

Table 4-2 Supported web browsers

Browser	Supported versions[a]
Mozilla Firefox	38
Microsoft internet Explorer	11
Google Chrome	43

a. IBM supports higher versions of the browsers as long as the vendors do not remove or disable functionality that the product relies upon. For browser levels that are higher than the versions that are certified with the product, customer support accepts usage-related and defect-related service requests. As with operating system and virtualization environments, if IBM Support cannot re-create the issue in our lab, IBM Support might ask the client to re-create the problem on a certified browser version to determine whether a product defect exists. Defects are not accepted for cosmetic differences between browsers or browser versions that do not affect the functional behavior of the product. If a problem is identified in the product, defects are accepted. If a problem is identified with the browser, IBM might investigate potential solutions or workarounds that the client can implement until a permanent solution becomes available.

2. Enter the Ethernet IP address of the frame on the URL line of the browser and press Enter.
3. If users are set up, enter your user name and password and click **Log in**. If this installation is new, log on to the TS4500 management GUI with the default user name and password (User Name: `admin` and Password: `admin`) to continue with the setup. After you log on with the default password, you must create a new password for the administrator account, which changes the default password.

> **Important:** Document the new administrator password because this password *cannot* be recovered if additional administrator users are not set up.

The TS4500 management GUI login window opens, as shown in Figure 4-3.

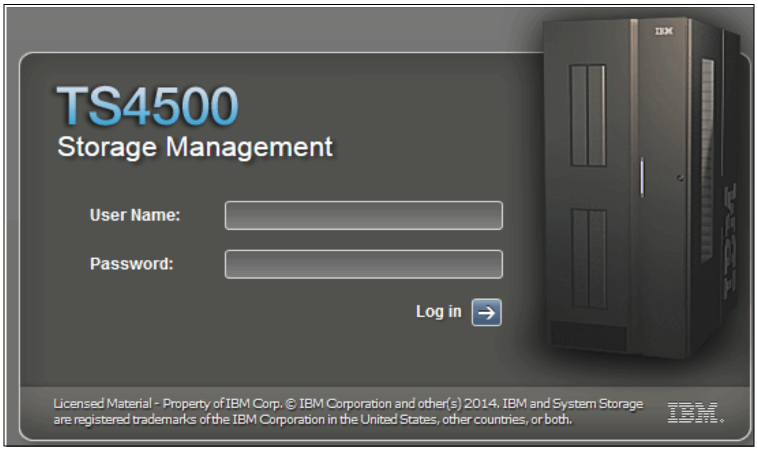

Figure 4-3 TS4500 login window

Setup wizard

If the TS4500 is a new installation, use the Initial Setup wizard, as shown in Figure 4-4. This wizard guides you through the basic configuration settings. The following settings are configured when you use the Initial Setup wizard. You can modify these settings from the Settings page of the management GUI.

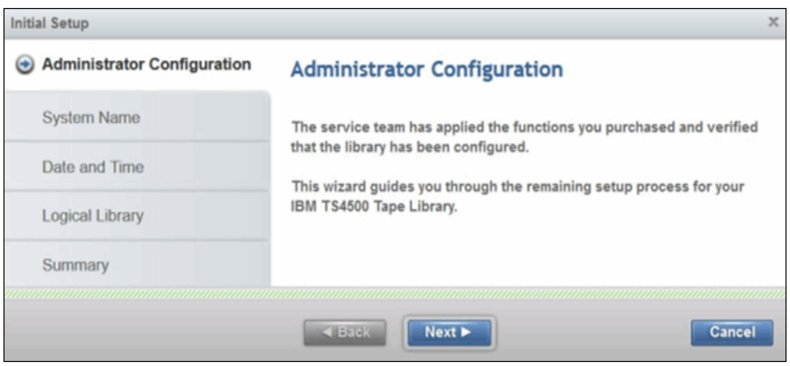

Figure 4-4 Setup wizard

Complete the following steps:

1. In the window, as shown in Figure 4-5, you can enter the name of your library.

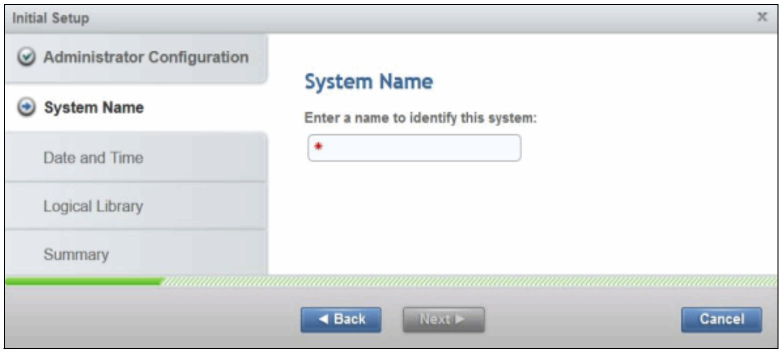

Figure 4-5 Setting the system name

2. In the next window, you can set the date and time for the library manually. Optionally, you can synchronize with a Network Time Protocol (NTP) server, as shown in Figure 4-6.

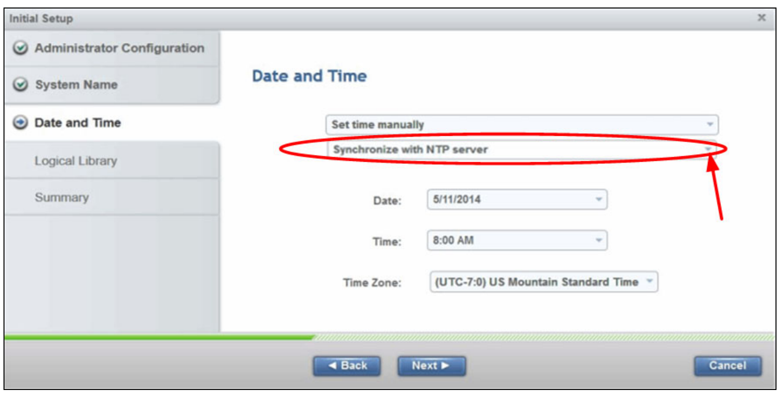

Figure 4-6 Date and time

3. Use the window that is shown in Figure 4-7 to configure a logical library, or you can skip this step and configure the logical library later.

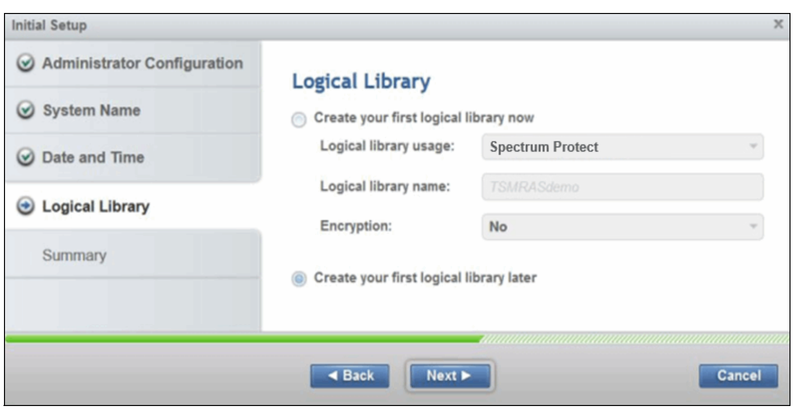

Figure 4-7 Logical libraries

4. The window that is shown in Figure 4-8 displays a summary of all of the settings that were configured in the previous steps. Select **Finish** to complete the wizard and go to the main menu.

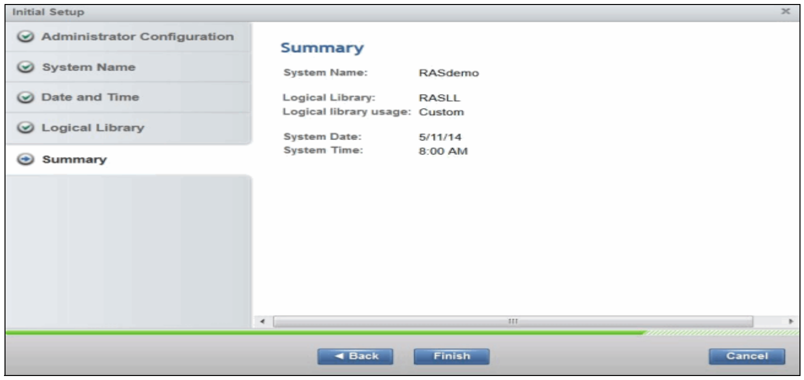

Figure 4-8 Summary

4.2.2 System summary display

After login the system summary window is shown, as in Figure 4-9.

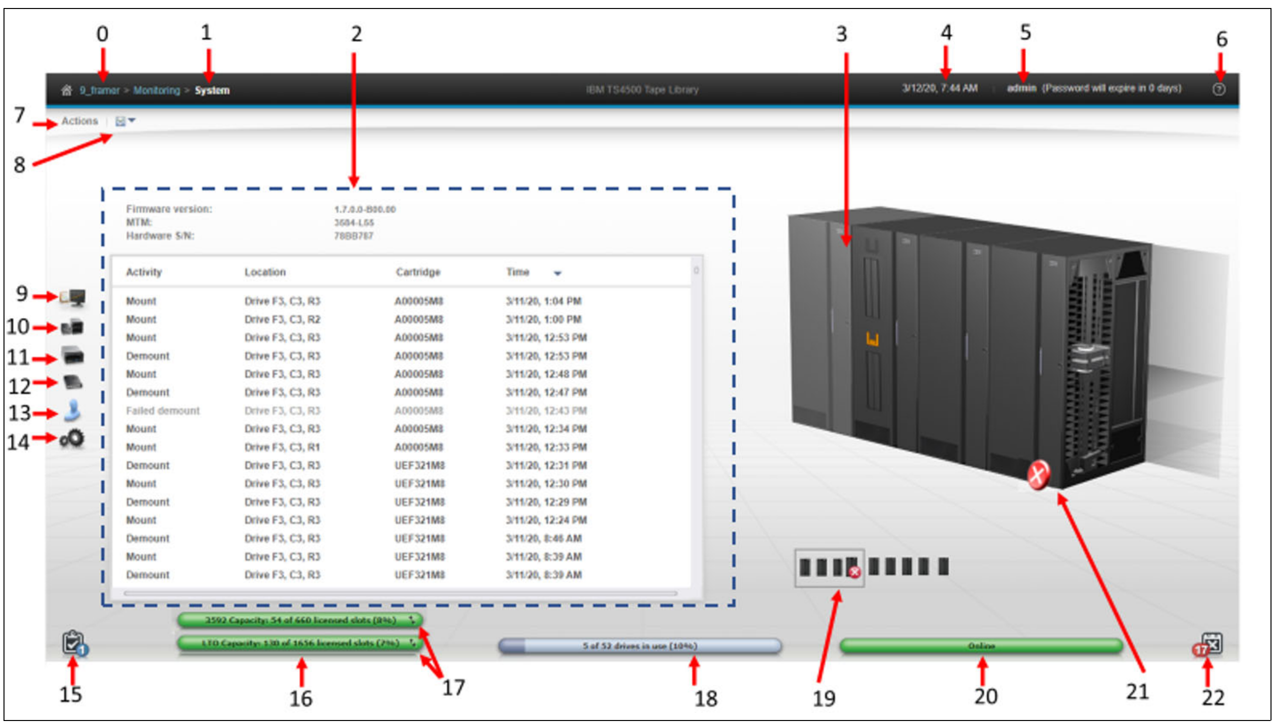

Figure 4-9 System summary display

For more information about the Management Interface (MI) and the available functions, use the question mark (**?**) icon or **Help**, which is number 4 in Figure 4-9.

Table 4-3 lists all of the fields that are shown in the window. The numbers in Table 4-3 refer to the numbers that are shown in Figure 4-9 on page 191.

Table 4-3 System display

Number	Field
0	TS4500 tape library name.
1	Current menu tree position.
2	The Activity Log displays the different activity types occurring in the library: Mounts, Unmounts, Imports, Exports, Moves, and Open/Close IO Stations.
3	View of physical library that is installed.
4	Library Date and Time
5	Logged-in user name and role.
6	Help.
7	Library hardware actions menu.
8	Export Data menu.
9	Monitoring menu.
10	Library menu.
11	Drives menu.
12	Cartridges menu.
13	Access menu.
14	Settings menu.
15	Tasks icon. The Tasks icon displays when tasks are running.
16[a]	The physical capacity pod displays the licensed capacity Linear Tape-Open (LTO) and (3592). This pod changes color to yellow or red, depending on whether the capacity exceeds the thresholds.
17	Switch the display to show the physical capacity per drive type by using the arrow on a mixed library.
18[a]	Drives in use status pod.
19[a]	All installed frames and status. This area is grouped into four frames. When this area is selected, the four frames show in the main window.
20[a]	Library health status pod. The color of the health status pod indicates the current state of the library by severity.
21	Frame health status pod for the selected frame. An icon is present in this field if an error or warning exists with a frame, tape drive, cartridge, or an accessor. Hover over the icon to see a list of the most important issues that cause this state. Click any issue to open the Events Page to see more information.
22	Library events icon. The library events icon displays if events are active.

a. Status PODs. Pods are always located at the bottom of the System page and show a quick view of capacity, drive utilization, and library health and changes color depending on the health or utilization settings.

4.3 Settings

Use the Settings menu as shown in Figure 4-10 to configure overall library options and settings. Each sub menu is described in the following sections.

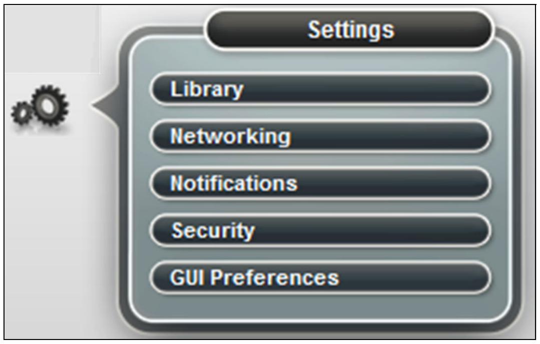

Figure 4-10 Settings

4.3.1 Library

Select Library on the Settings menu to access the Library page to set the date and time, configure the advanced options, and work with licensed functions.

Date and Time option

Use the Date and Time option to set the library date and time. You can choose whether to set the date and time manually (see Figure 4-11) or to synchronize with the NTP server (see Figure 4-12 on page 194). If you select to use the NTP server, you can optionally insert a primary and secondary NTP server and test the connection from this display.

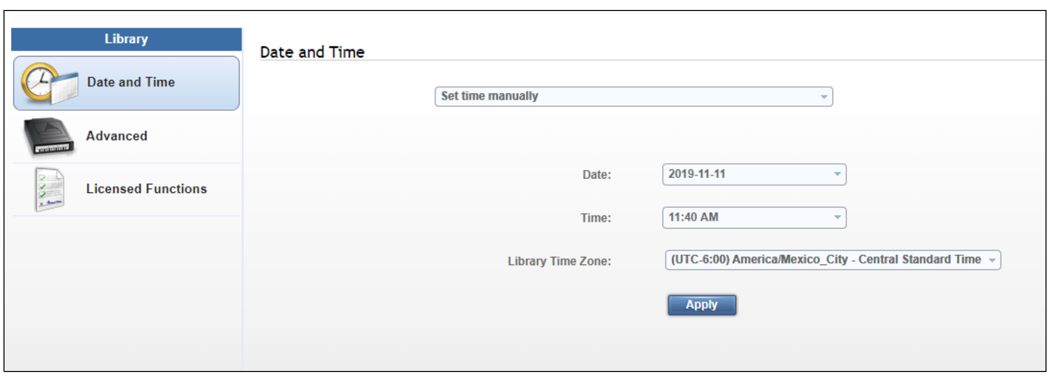

Figure 4-11 Set time manually

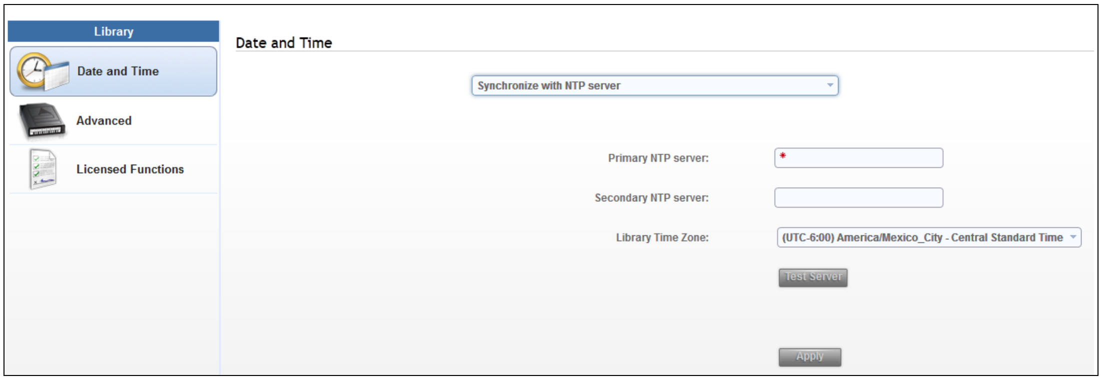

Figure 4-12 Synchronize with NTP Server

Advanced

Use the Advanced page to manage expired cleaning cartridges, and REST over SCSI as shown in Figure 4-13. If you have dual accessors the Advanced page will additionally show elastic capacity and active accessor options, as shown in Figure 4-14.

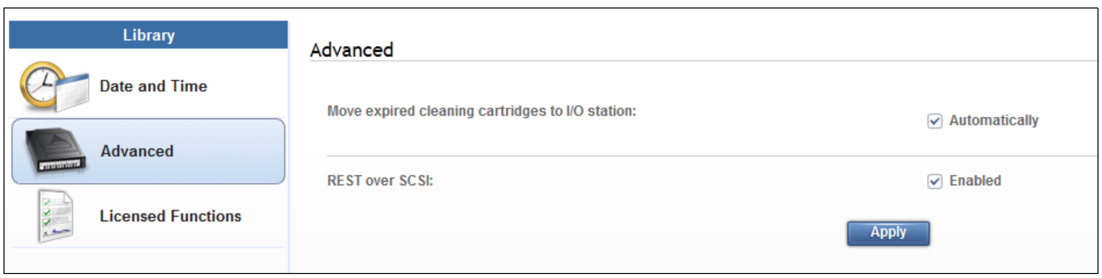

Figure 4-13 With single accessor

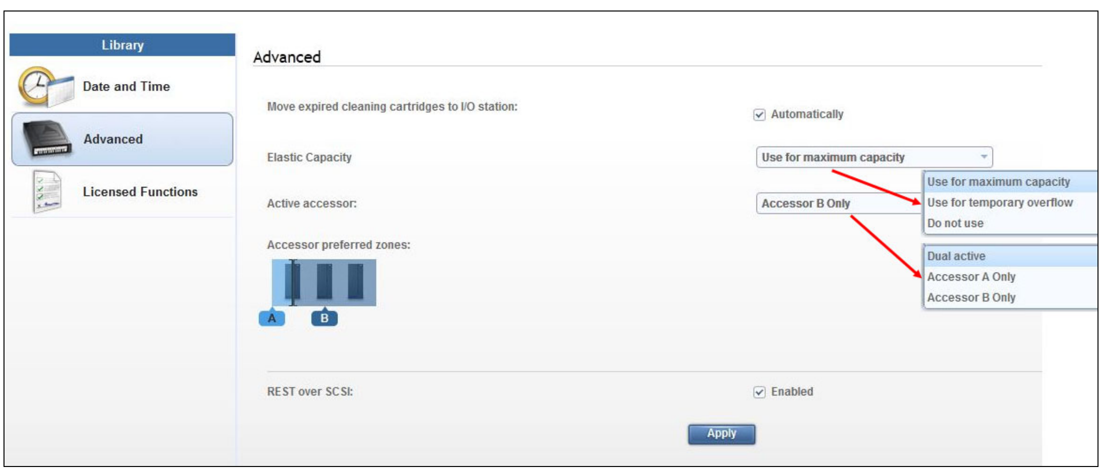

Figure 4-14 With dual accessors

Moving expired cleaning cartridges to I/O station

The library uses the cleaning cartridge to automatically clean tape drives as needed to maintain the efficiency of the tape drives. This option allows the TS4500 to automatically move expired cleaning cartridges to the I/O station.

Each cleaning cartridge can be used 50 times. The usage count is stored internally in the cartridge memory chip inside each cleaning cartridge, which prevents an expired cleaning cartridge from being reused.

Elastic Capacity option

If the library has the High Availability (HA) feature (dual accessors), use this Elastic Capacity option to manage the way cartridges are stored in the accessor service areas. The following settings are available:

- Use for maximum capacity: The least recently used cartridges are moved to an accessor service area when the non-service area is 98% full.
- Use for temporary overflow: The I/O station cartridges are imported to an accessor service area when the non-service area is 100% full.
- Do not use: Cartridges are never moved to an accessor service area.

The accessor service area consists of the storage slots that only one accessor can manage. For more information, see “Elastic Capacity option” on page 24.

Active accessor

Use this Active Accessor option to enable and disable accessors, as required. If a service issue occurs with one accessor, you can select the accessor that works to service the whole library and disable the accessor that requires service. Figure 4-15 shows the display for setting either accessor A or accessor B as active and shows a representation of the available slots with a single accessor set.

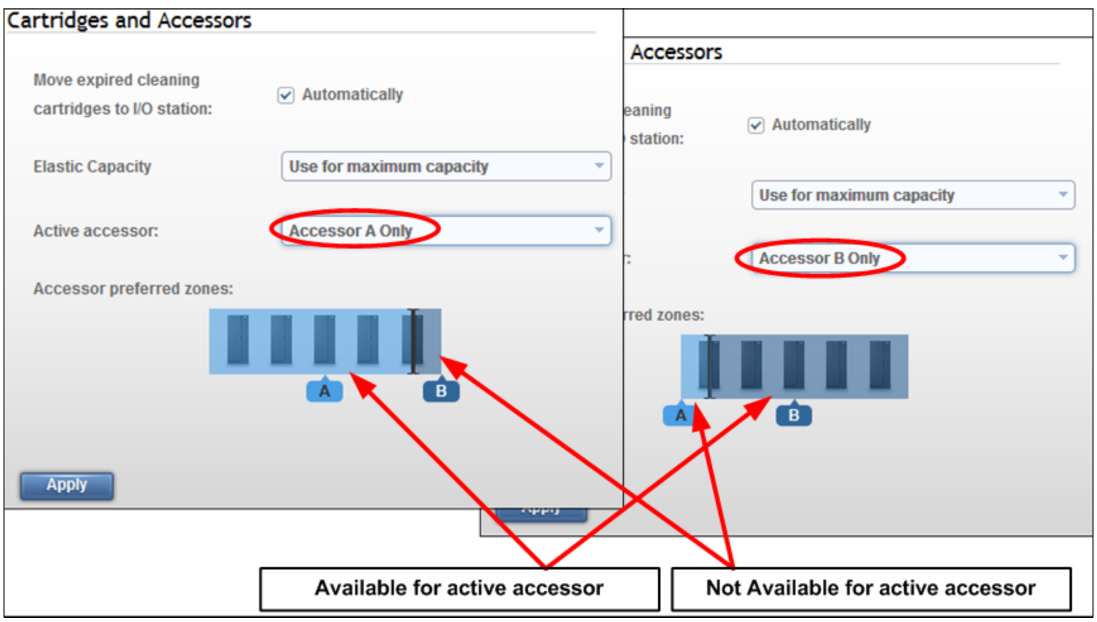

Figure 4-15   Single active accessor

Accessor preferred zones

The Accessor preferred zones are read-only graphics, which display the preferred zone for each accessor, including the accessor service areas if Elastic Capacity is enabled.

Modifying the accessor zone configuration is done using the TS4500 **`setAccessorZones`** CLI command. This command can also be used to inactivate an accessor. This command is described in “viewAccessorZones” on page 371.

REST over SCSI

Use this option to enable (default) or disable REST over SCSI. REST over SCSI is described in detail in Chapter 6.

Licensed functions

Licensed functions enable extended library capabilities that are available only to users who purchase and enter a license key code for that particular capability. All of the available extended capabilities are listed in the table on the Licensed Functions page. A green check mark in the Licensed column indicates that the license is installed.

You are required to enter the license key in the license key code field. The license key file is generated by manufacturing, and it is a unique number that is based on the serial number of the library. Select the frame number and enter the license key, as shown in Figure 4-16.

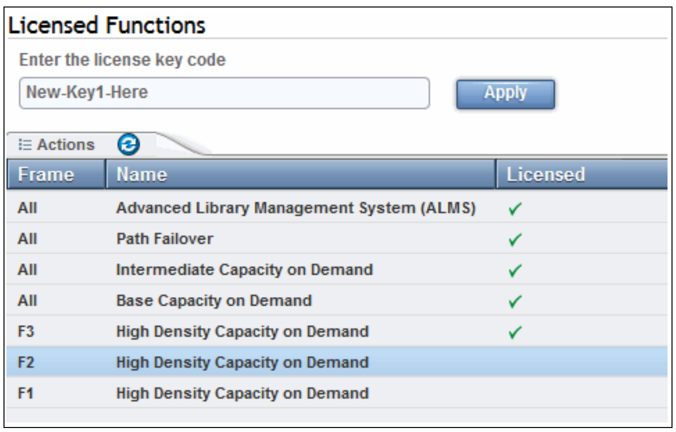

Figure 4-16 Licensed function

> **Note:** Advanced Library Management System (ALMS), which comes standard on the TS4500, always shows as installed.

The following features are the extended library capabilities and they become available when the license key is uploaded:

- *Path failover* creates redundancy in the path from the application to the intended target (the library accessor or the drive mechanism).
- Intermediate high-density capacity on demand (HD CoD) increases storage from the entry capacity to the intermediate capacity. The Intermediate CoD feature (FC 1643) adds 100 slots, increasing the usable capacity of the L25 and L55 frames to 200 slots.
- Base capacity on demand (CoD) increases storage from the intermediate capacity to the base capacity. The Base CoD feature (FC 1644) adds 200 slots, increasing the usable capacity of the L25 and L55 frames to 400 slots.
- LTO transparent encryption is required to enable encryption on LTO tape drives if you use library-managed encryption (LME).
- High-density capacity on demand (HD CoD) increases storage to use all of the tiers of an HD frame. The initial capacity of the Dx5 frames is 500 slots. The initial capacity of the S25 frame is 600 slots. The initial capacity of the S55 frame is 660 slots. The HD CoD features can add 50 - 660 more slots anywhere, depending on the frame position and configuration. This license is required to be installed on each HD frame.

- It is possible to remove a licensed function by using right-click a licensed function and selecting remove license key as shown in Figure 4-17.

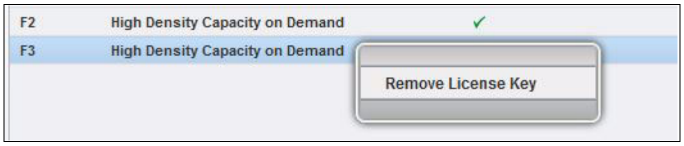

Figure 4-17 Remove License Key

4.3.2 Networking

Select Networking on the Settings menu to access the Networking page to set up management Ethernet ports and iSCSI Ethernet ports on the TS1160 model 60E and TS1155 model 55E.

Management Ethernet Ports

Use the Management Ethernet Ports menu to display and configure the ports on each library controller card (LCC). Each port can enable or disable a particular protocol. The LCC that is specified for each port is the LCC to which that network connects when local hardware communicates with remote hardware.

By selecting the **Network** option, the Ethernet Ports window opens, as shown in Figure 4-18, which displays the current settings of all ports and displays their status.

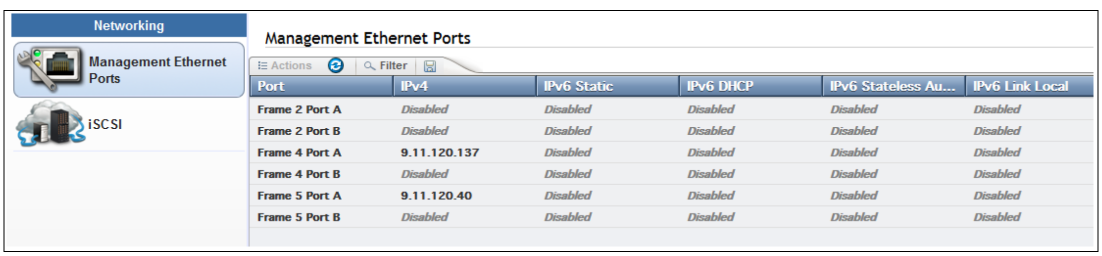

Figure 4-18 Ethernet Ports window

Use this window to change the Internet Protocol (IP) address that you use to access the management GUI. Individual frames and ports can be configured to use IPv4, IPv6, or both types of IP addresses.

Use the Actions menu to modify an Ethernet port or display its properties, as shown in Figure 4-19.

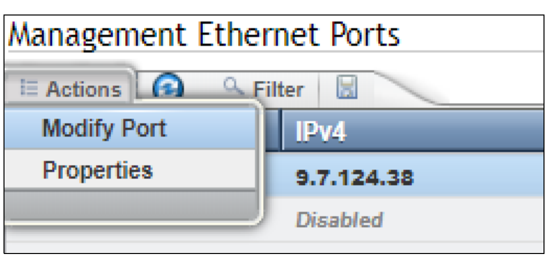

Figure 4-19 Ethernet ports

Modifying a port

To modify a port, highlight the port that you want to modify and select **Actions** → **Modify Port**. The Modify Frame window opens, as shown in Figure 4-20.

Figure 4-20 Modify Ethernet port settings

Use this window to change the IP address that you use to access the management GUI. You can configure individual frames, and ports can be configured to use IPv4, IPv6, or both types of IP addresses.

The following fields are available:

- The IPV4 assignment can be set to static, DHCP, or disabled:
  - Static or fixed IP addresses are manually assigned to each system by an administrator.
  - For Dynamic Host Configuration Protocol (DHCP), set this field to DHCP if you use a DHCP server, which automatically configures the IP address and network parameters.
- The IPV4 address, subnet, and gateway fields are used to set up a static IP configuration.
- The IPV4 primary and secondary Domain Name System (DNS) fields are used to set up the IP address of a DNS server, if required.
- The IPV6 assignment can be set to disabled, DHCP, stateless auto configure, or static:
  - The static or fixed IP addresses are manually assigned to each system by an administrator.
  - Set this field to DHCP if you use a DHCP server, which automatically configures the IP address and network parameters.

– Stateless auto configure allows a host to generate its own addresses. It uses a combination of the router prefix (identifies the subnet that is associated with a link) and a host-generated interface identifier (uniquely identifies an interface on a subnet).

- The IPV6 address, subnet, and gateway fields are used to set up a static IP configuration.
- The IPV6 primary and secondary DNS fields are used to set up the IP address of a DNS server, if required.

A link local address is an IP address that is intended only for communications within the segment of a local network. The link local address is not configurable. IPv6 is enabled per interface, and the IPv6 link local address is assigned to the interface where IPv6 is enabled.

Properties

This option displays the current configuration and link status of the selected port, as shown in Figure 4-21.

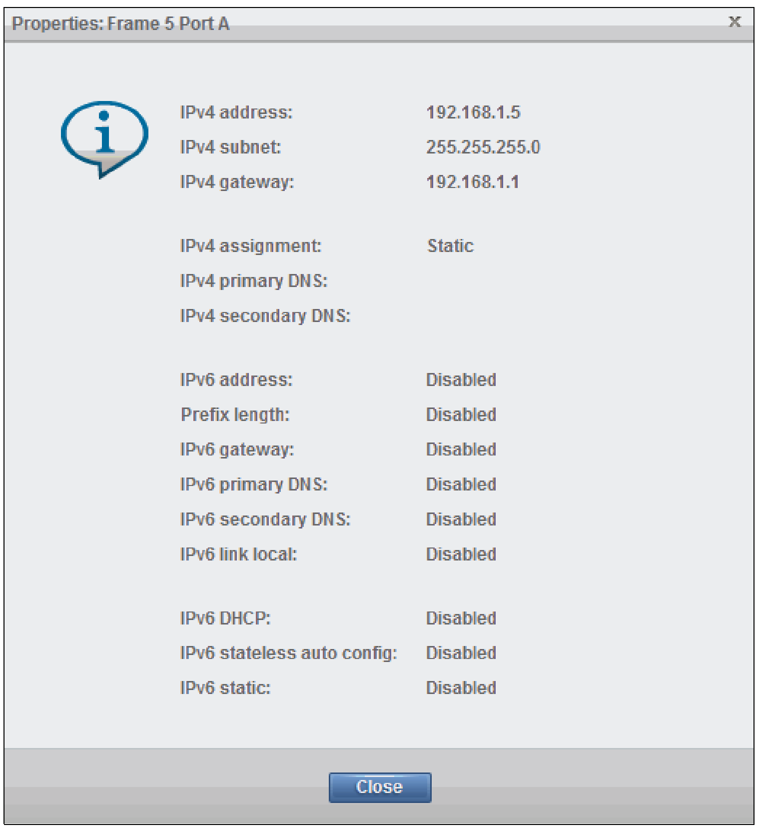

Figure 4-21 Ethernet port properties

iSCSI

This option provides for library-wide security setting for the TS1160 model 60E and the TS1155 model 55E iSCSI Ethernet ports.

Figure 4-22 shows the security setup options that are available for iSCSI.

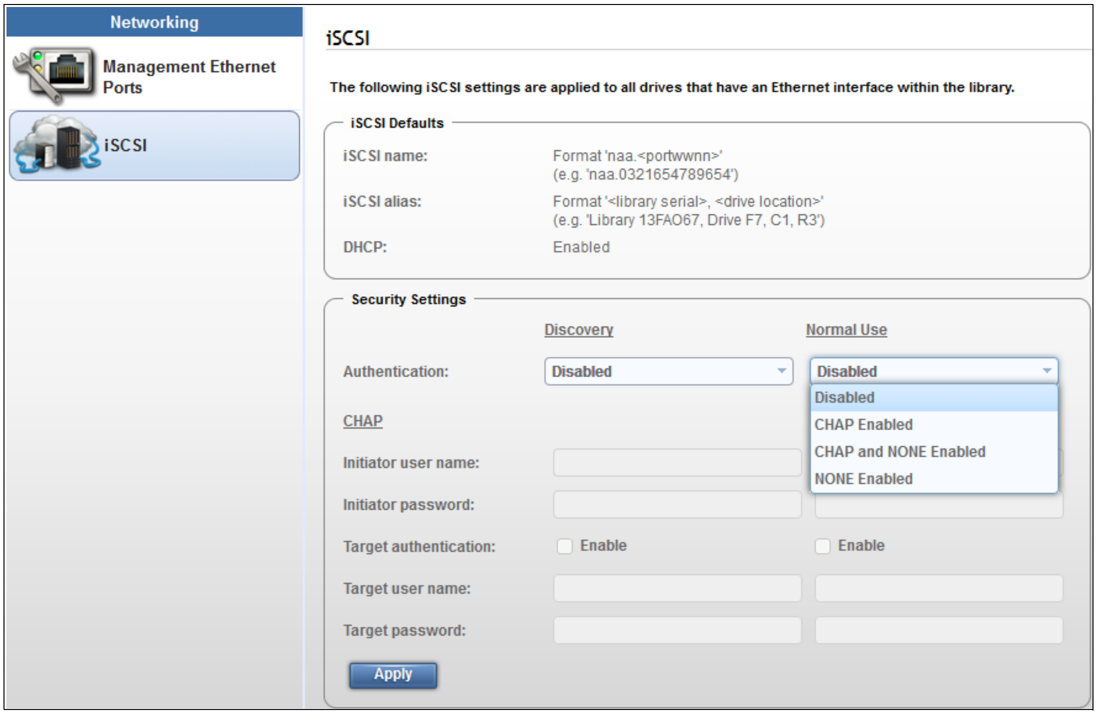

Figure 4-22 iSCSI library wide security settings

The following are possible authentication settings for discovery and normal use:

- Disabled, which is the default
- CHAP enabled
- CHAP and NONE enabled
- NONE enabled

When Challenge-Handshake Authentication Protocol (CHAP) support is enabled, hosts are securely authenticated by the system. This increases overall system security by verifying that only authenticated parties are involved in host-storage interactions.

CHAP is an authentication process of an iSCSI initiator by a target through comparing a secret hash that the initiator submits with a computed hash of that initiator's secret, which is stored on the target.

Note: Target authentication is disabled by default. Target user name and password are required if using CHAP and Target authentication is enabled.

4.3.3 Notifications

Select Notifications on the Settings menu to access the Notifications page to configure the sender information, recipient information, and library information for the various TS4500 tape library notifications. The Notifications page contains several options, as shown in Figure 4-23.

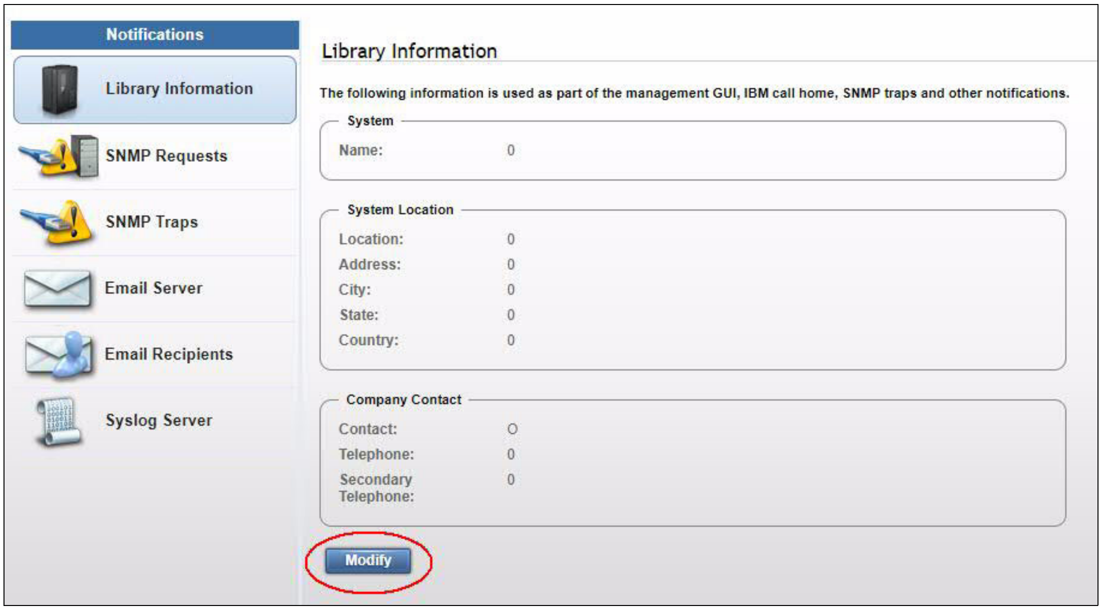

Figure 4-23 Notifications page

Library Information

The Library Information page displays the configured system name, system location, and company contact information for the TS4500 tape library. The library information is sent in Simple Network Management Protocol (SNMP), email, syslog, and Call Home notifications.

To modify all fields, complete the following steps:

1. Click **Modify** on the Library information window, as shown in Figure 4-24.
2. The Modify option opens the window that is shown in Figure 4-24. On the Modify Library information window, you can set all of the system information and configure the library name (which can be changed concurrently).

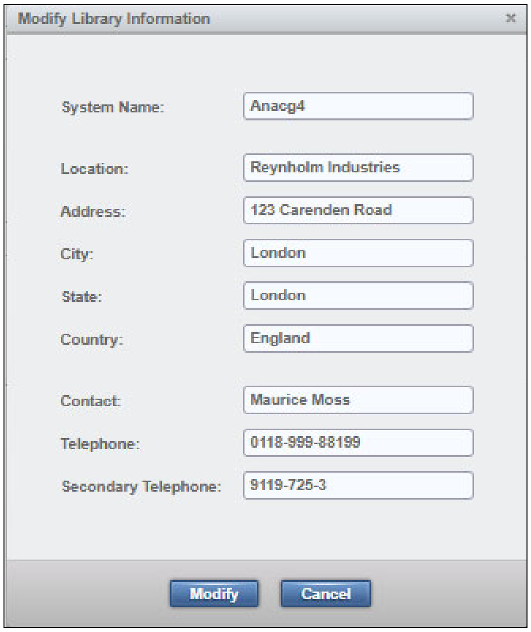

Figure 4-24 Modify Library information window

3. Use this window to set the TS4500 system name, system location, and company contact details. The system name that is shown in the library information profile references the system name that is set during the initial configuration of the library. The system name appears in the navigation tree and notifications.
4. After you complete all of the settings, click **OK** to apply the changes.

SNMP Requests options

To set the SNMP Requests options, complete the following steps:

1. Select **Notifications** on the Settings page, as shown in Figure 4-10 on page 193, to configure how the TS4500 tape library sends SNMP traps and requests. SNMP notifications include SNMP traps and SNMP requests.

   SNMP traps enable the tape library to send its profile to the SNMP server by way of an unsolicited SNMP message.

2. Under the Notifications menu, select **SNMP Requests** to display the current settings, as shown in Figure 4-25.

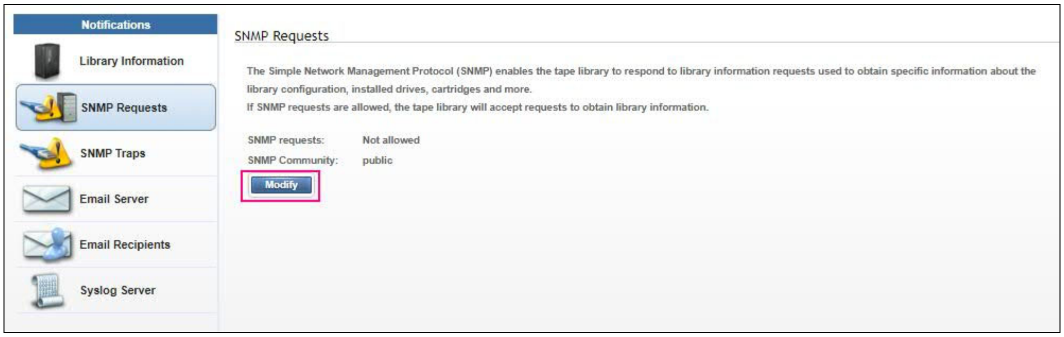

Figure 4-25 SNMP Requests

3. To change the setting, select **Modify**, as shown in Figure 4-25, which opens the Modify SNMP Settings window, as shown in Figure 4-26.

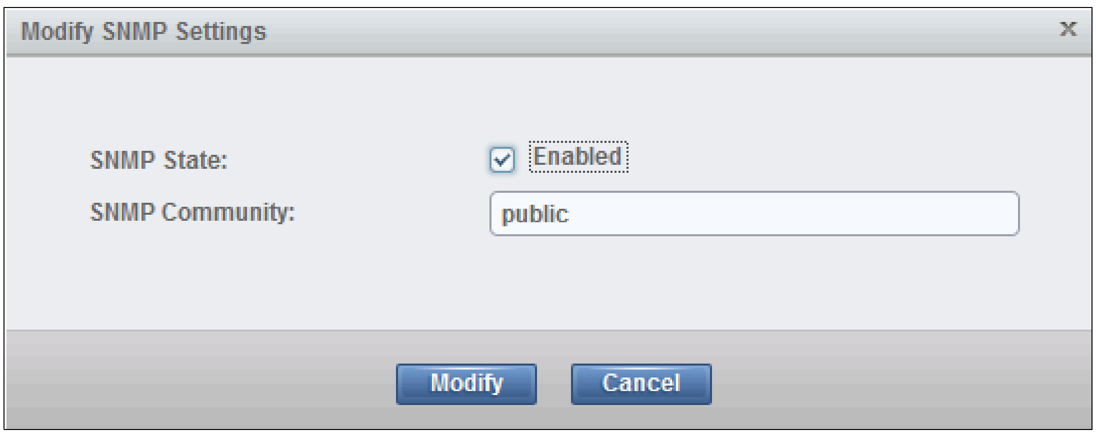

Figure 4-26 Modify SNMP Settings window

4. Select the SNMP state (Disabled or Enabled). When the SNMP state is Not allowed, the TS4500 tape library does not send messages to the network-attached devices for conditions that warrant administrative attention. When the SNMP status is Allowed, the library sends traps to automatically notify an administrator if an issue arises.

> **Note:** Disabling SNMP does not delete any destinations that were set up.

SNMP is a networking protocol that, when it is enabled, allows the TS4500 tape library to automatically gather information about alerts and status. The system then transmits this information to other entities, such as an SNMP monitoring server, in the network. The gathered information is called an *SNMP trap*.

SNMP traps enable the TS4500 tape library to send its profile to the SNMP server by way of an unsolicited SNMP message. If an issue arises with the library, network, or any port in the network, the tape library responds with an information profile to the SNMP server. The nature of the profile depends on the type of issue that arose.

The *SNMP community* is the name of the class of users that can access the statistics of network-attached devices. The trap community name is sent with a trap. For information about the SNMP trap community name, see the documentation for your monitoring station.

To view or change the community name that is associated with the TS4500 tape library, click **Modify** in the Modify SNMP Settings page (see Figure 4-26 on page 203).

By default, the TS4500 tape library SNMP community is set to `public`. During the initial system configuration, the administrator can change the community name and customize the access settings for each community name.

SNMP Traps

Use the SNMP Traps window to configure SNMP destination servers, send test SNMP traps, and download the SNMP Management Information Base (MIB) file. You can also specify the type of messages (errors, warnings, or informational messages) that each SNMP destination server receives:

1. From the SNMP Destinations window, select the **Create Destination** tab to display the Add Destination window, as shown in Figure 4-27.

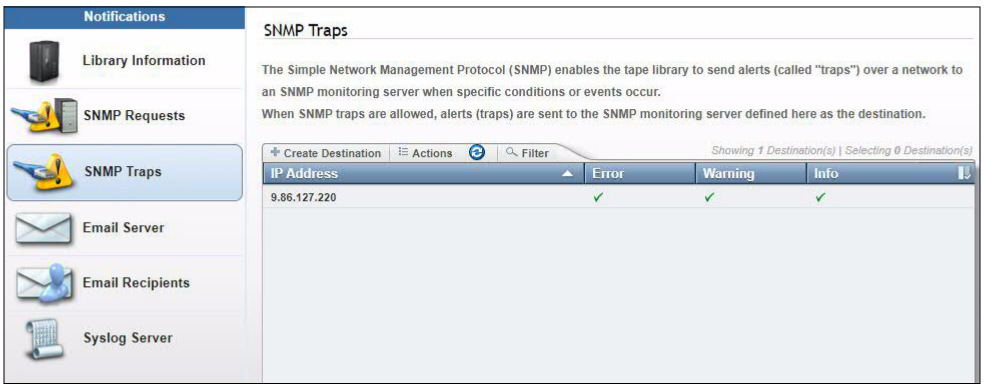

Figure 4-27 SNMP Traps window

2. Enter the destination SNMP server IP address and port number that are used, with the types of messages to send to the server.
3. Select **Add** to add the server and continue to create new servers if you require multiple servers with different roles.

   You can select errors, warnings, and informational messages to send to the SNMP server.
4. From the SNMP Traps window, select **Actions** to modify or delete an existing destination server, send a test trap, or download the SNMP MIB file. Figure 4-28 shows these options.

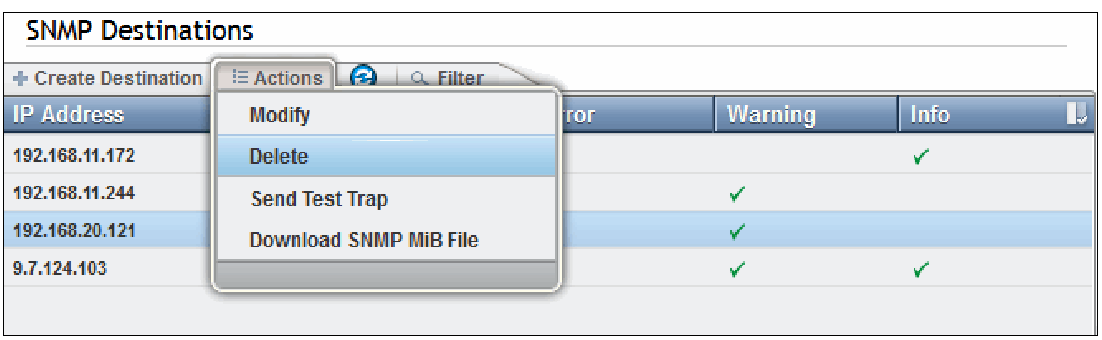

Figure 4-28 SNMP Traps actions

The SNMP MIB file is used to interpret SNMP traps. When the GUI receives an SNMP trap, it compiles the SNMP trap into human-readable form so that you can gather information about the error.

The many components in a network are made by various manufacturers, and each component has unique properties and definitions. Data that is sent from the TS4500 tape library to any of the devices must be translated to a protocol that is understood by the device. An *MIB file* is a database that contains the properties and definitions of each network-attached device. The MIB file receives and translates the data when the TS4500 tape library sends the request for information.

SNMP query configuration

The TS4500 tape library stores its major configuration components in a standard Management Information Base (MIB) file. You can use an SNMP `GetRequest` to query the library and use the configuration MIB to translate the fields that gather configuration data for the TS4500 tape library and other IBM library types by using the SNMP query feature. The SNMP query feature is an easy way to gather configuration data from both local libraries and libraries that are in different geographic locations.

The following list shows all of the configuration data options that you can access for the frame or system configuration (frame or module):

- Logical Library
- Drive Configuration
- Library Configuration
- VPD Note Card
- Call Home Configuration
- SNMP
- Key Manager
- Drive Encryption
- SMTP configuration
- Time Configuration
- User Roles
- Role Access Level
- Lightweight Directory Access Protocol (LDAP) Configuration
- Ethernet Configuration

To download the configuration MIB file, select **SNMP Traps**, as shown in Figure 4-28 on page 204. Then, with the IP address highlighted, select **Actions** → **Download SNMP MiB File**.

Email Server

Email is a method, other than SNMP traps, to send information to users who need information about events that occur in the network. If your Simple Mail Transfer Protocol (SMTP) server requires authentication information before it accepts email notifications, you can set the authentication information, including the SMTP server port number, by using the Modify Email Settings page:

1. To access the Modify Email Settings page, click **Modify** on the Email Server page.
2. Enable **SMTP authentication**, as shown in Figure 4-29 on page 206.

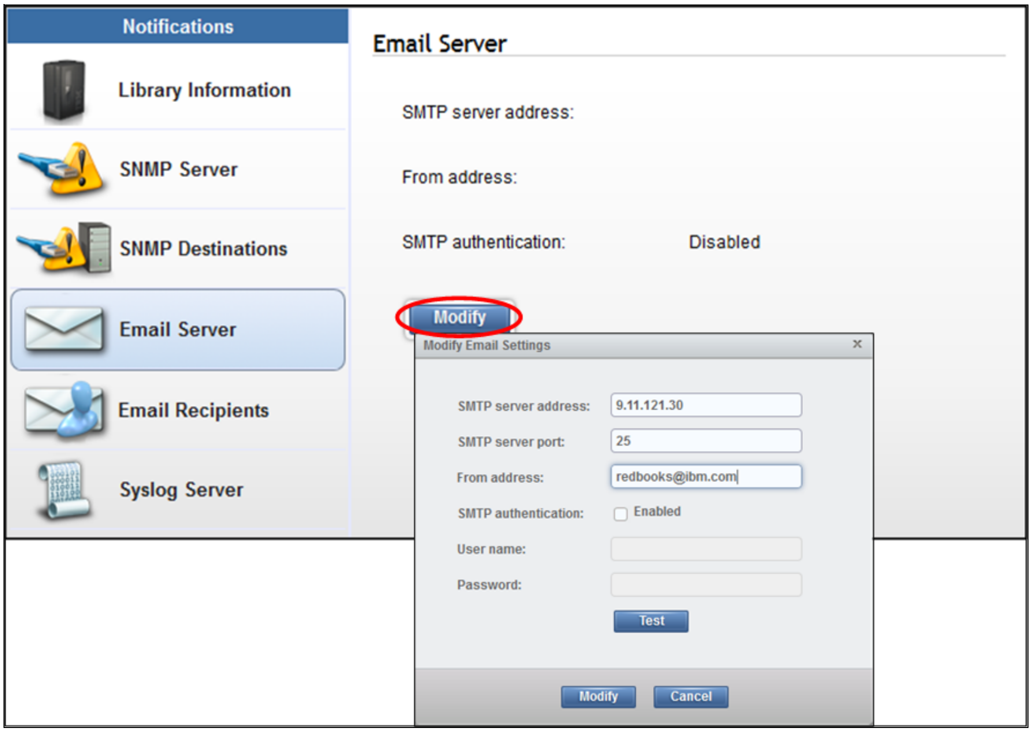

Figure 4-29 Email Server notifications page

3. After you complete all fields, click **Test** on the Modify Email Settings page. If the test is successful, click **Modify** to save the settings.

SMTP dictates that every email message must specify the email address of the sender. This sender address must be a valid address for the following reasons:

- Many SMTP gateways require a valid sender address as a security measure to prevent unauthorized usage of the SMTP server. Otherwise, the SMTP gateway does not forward the email. Often, this sender address must be limited to a specific domain.
- The sender's address is used as the destination for error messages, such as an incorrect email address and a full email mailbox, that are generated by the SMTP gateways.

Note: The default SMTP server port value is set to 25, and can be changed to a secure port such as 465. Many email serves no longer permit port 25 so check with your email service provider or administrator for correct port to use.

Email Recipients

Use the Email Recipients page to add email addresses for recipients of the error notifications, warning notifications, and informational notifications:

1. Select **Email Recipients**.
2. Select the **Create Recipient** tab, as shown in Figure 4-30.

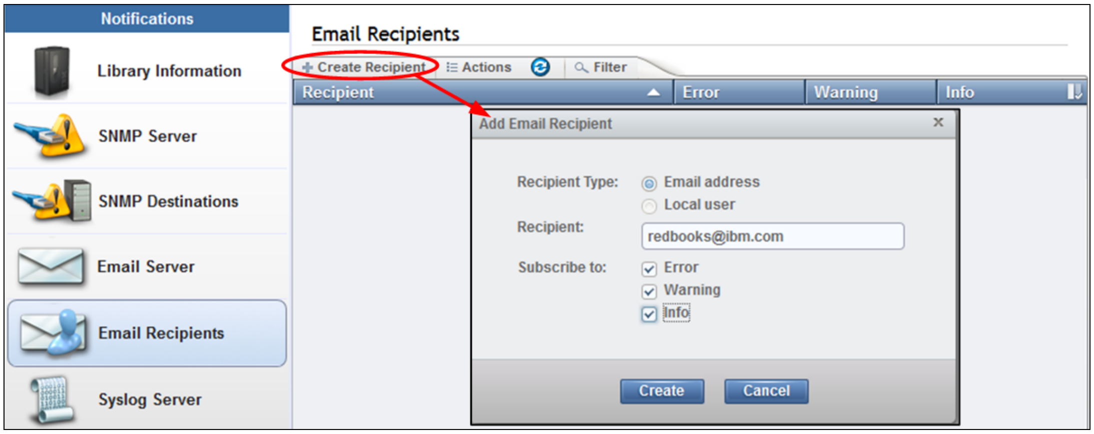

Figure 4-30 Email Recipients option

3. You can add a number of recipients, based on their email addresses and local user names. Specify the email address to which you want to send the events, or you can optionally select a user if an email address was specified for that user when that user was defined.
4. You can also select the type of alert that a specific user can receive. Select **Create** to create an email recipient.
5. After a recipient is created, you can then modify the user, delete the user, or send a test email to the user from the Actions menu, as shown in Figure 4-31.

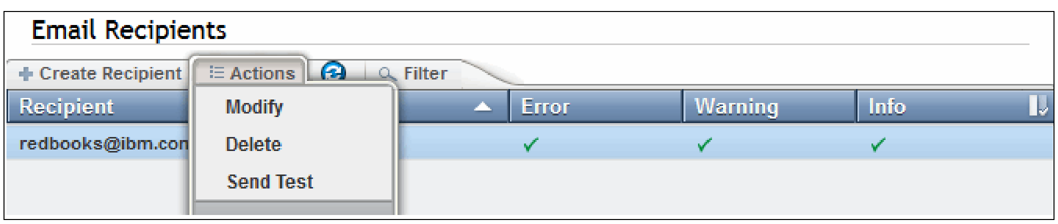

Figure 4-31 Email Recipients Actions menu

System log (Syslog) server

Use the **Syslog Server** option on the Notifications menu to set up a destination syslog server, as shown in Figure 4-32:

1. Click the **Create Recipient** tab.
2. Configure the IP address of the syslog server, the server port number and subscribe to the types of events to send to the syslog server. Click **Create**.

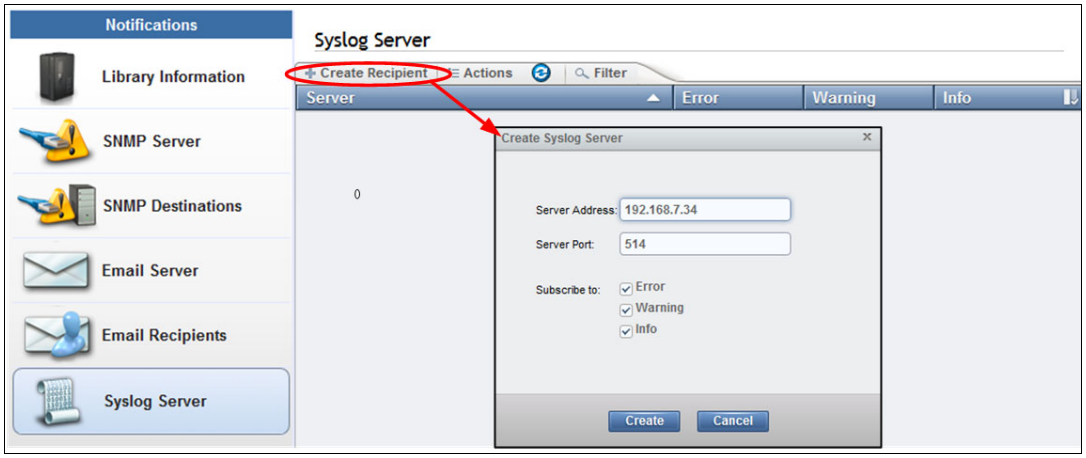

Figure 4-32 Syslog Server page

When system events occur, the TS4500 tape library creates a log of these events. You can configure the TS4500 tape library to send syslog notifications and send a notification of the event to the syslog server. The syslog server keeps its own log of system events. (The syslog server is a client-provided server.)

3. After the syslog server is set, you can use the **Actions** menu, as shown in Figure 4-33. You can modify, delete, or send a test to the IP address of a syslog server.

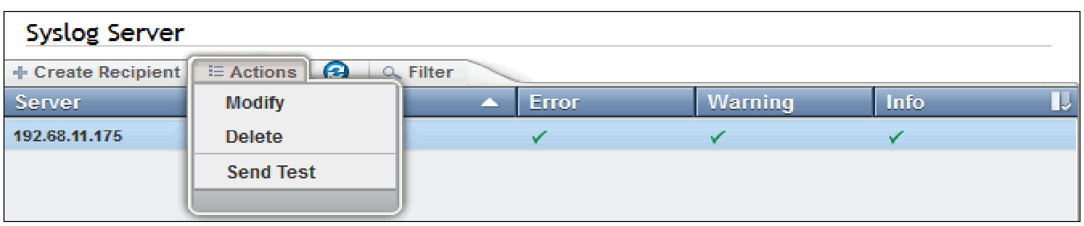

Figure 4-33 Syslog Server Notifications Actions menu

4.3.4 Security

Select **Security** in the Settings menu to access the Security page for the following functions:

- Enable and configure remote authentication (LDAP)
- Disable remote authentication
- Set access rules
- Enable or disable Secure Sockets Layer (SSL) for GUI and CLI communications with the library
- Manage encryption

Remote authentication

With remote authentication, security tasks are centralized and user management can be performed from a single interface, without logging in to the TS4500 tape library. This capability was referred to as the *Storage Authentication Service* (SAS) with the TS3500 tape library. The default status for remote authentication is `Disabled`, as shown in Figure 4-34.

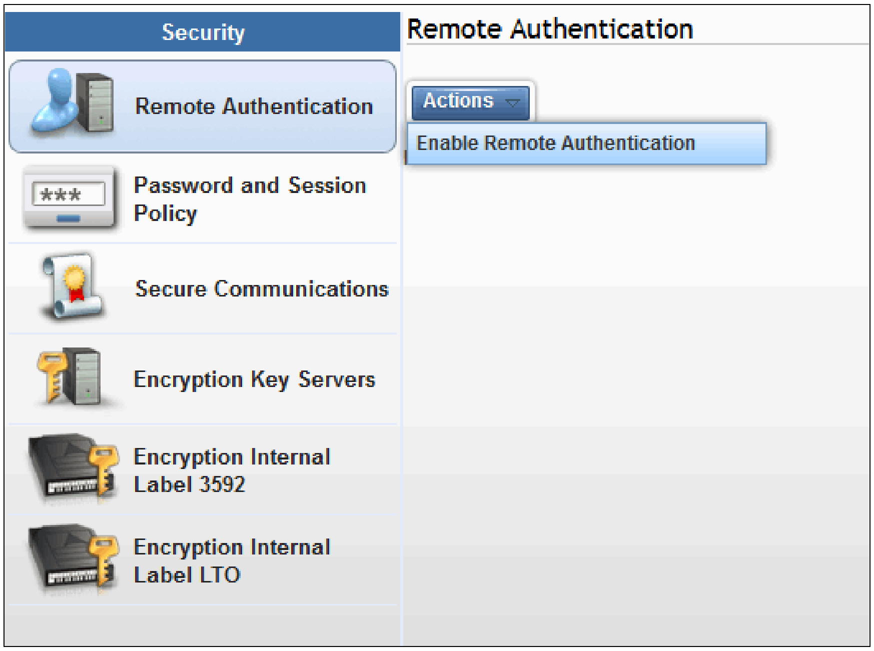

Figure 4-34 Security Remote Authentication page

In this section, we discuss LDAP as the main remote authentication service. The same authentication applies to IBM RACF® unless otherwise noted. For more information about the differences in terms of RACF configurations, see “Resource Access Control Facility” on page 216.

Disabled status

With disabled status for remote authentication, users are defined on the local library and the library uses local authentication to manage the access. For more information about user set up, see 4.4.1, “Users” on page 227.

Enabled status

With remote authentication enabled, the users are defined on the LDAP-Server. All authentication requests are passed to this remote authentication server that verifies the user’s name and password. Then, the LDAP-Server returns a list of groups that is associated with this user. This list of groups is then mapped to a customized role within the TS4500.

To enable remote authentication, the following conditions must be met:

- A group is created and users are added to that group on your remote authentication server.
- On the TS4500 tape library, a custom role is created with a name that matches the group name on the remote authentication server. We support characters [a-z][A-Z][0-9] and “-” and “_” with a maximum length of 50 characters. Roles with spaces are not supported.
- Each custom role can have different access permissions.
- The primary LDAP repository Uniform Resource Identifier (URI) is required.

- The secondary LDAP repository URI is optional.
- The LDAP Transport Layer Security (TLS) certificate is optional.
- If you use Kerberos, you require the realm, key distribution center (KDC), and the domain mapping.

Lightweight Directory Access Protocol

LDAP is an open protocol that uses TCP/IP to provide access to directories that support an X.500 model. For example, LDAP can be used to locate people, organizations, and other resources in an internet or intranet directory. In our scenario, users are mapped to groups and roles.

The basic concept is that a user is authenticated on the LDAP-Server and the TS4500 queries for all the groups in which the user belongs. These groups are mapped to customized roles within the TS4500. The user is assigned with the access permissions that are set within the customized role.

If a user is a member of multiple groups that correspond to TS4500 roles, the user inherits the access level of the first role match that is found in alphabetical order. For example, if a user belongs to two groups that have corresponding TS4500 custom roles (fore example, `tech_admin` with Administrator access and `tech_service` with Service access), the user is has Administrator-level access.

Kerberos

Kerberos is a network authentication protocol that is based on symmetric key cryptography. Kerberos assigns a unique key, which is called a *ticket*, to each user who logs on to the network. The ticket is embedded in messages that are sent over the network. The receiver of a message uses the ticket to authenticate the sender.

Kerberos settings include the following characteristics:

- The Kerberos realm is generally the same as your company’s domain name. For example, if your company’s domain name is `example.com`, your Kerberos realm is `EXAMPLE.COM`.
- The Key Distribution Center (KDC) (AD server) is the Key Distribution Center server. A KDC server generally has a prefix of “`Kerberos`” followed by your Kerberos realm, a colon, and the port number of the Kerberos server. (The port number of the Kerberos server is `88` for the TS4500 tape library.) Therefore, if your company’s domain name is `example.com`, a conventional name for your KDC server is `kerberos.example.com:88`.
- Domain Name Service must be configured to enable Kerberos on the library and it is highly recommended to configure the library to use the same NTP server that the KDC is configured to use.

Custom roles

With remote authentication, access control is managed by the remote authentication server.

Each custom role has unique access permissions. For more information about custom roles, see 4.4.2, “Roles” on page 234.

Enable remote authentication

After you complete all required preparation and configuration, select **Actions** to enable remote authentication, as shown in Figure 4-34 on page 209.

The Remote authentication window opens. The first page lists the prerequisite actions before you can start the wizard to enable remote authentication, as shown in Figure 4-35 on page 211.

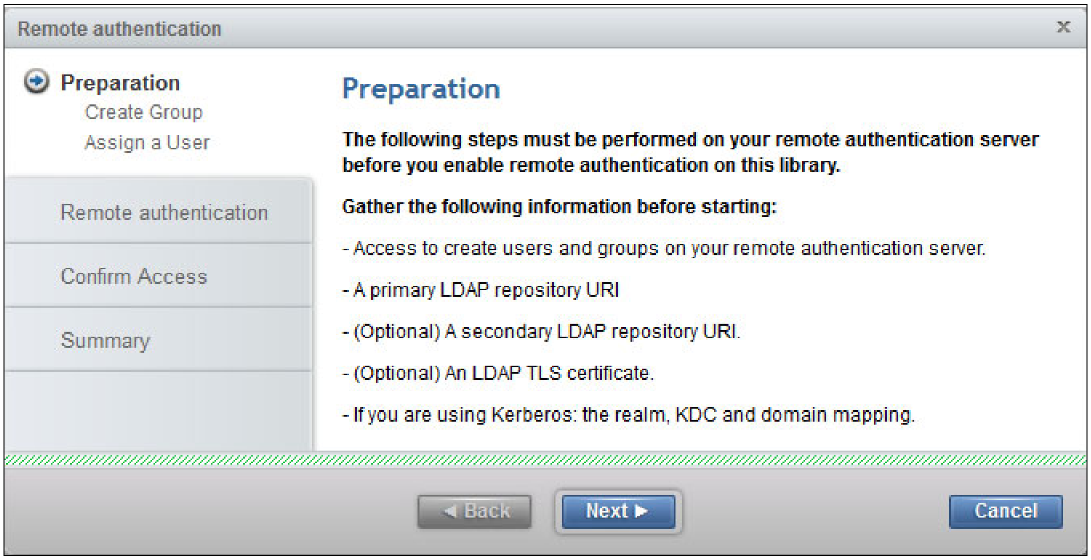

Figure 4-35 Remote authentication Preparation window

The next two steps require you to confirm the custom Roles (on the TS4500) to match the groups on the LDAP-Server and to add users to this group. Consider the following points:

- The Create Group page shows you the created custom Roles. At least one role must match one group on the LDAP Server (only roles with Administrator access are displayed).
- The Assign a User page reminds you to add a user to this group on the LDAP-Server. The maximum supported length of the username is 50 characters.

If you intend to user Kerberos, select **Kerberos**; otherwise, select **LDAP** (see Figure 4-36).

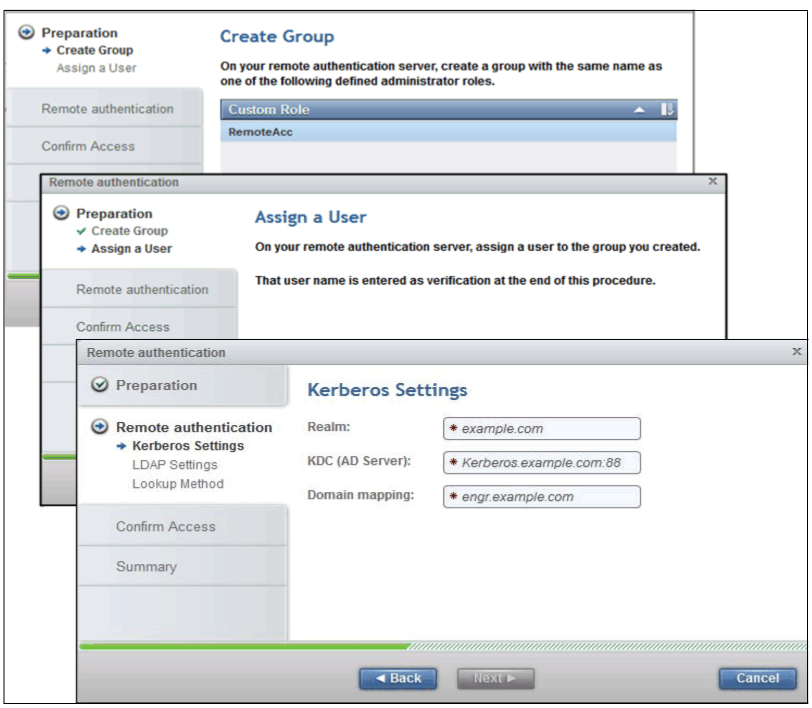

Figure 4-36 Remote authentication preparation

Kerberos

If you selected the Kerberos authentication method, the Kerberos Settings window opens (see Figure 4-37).

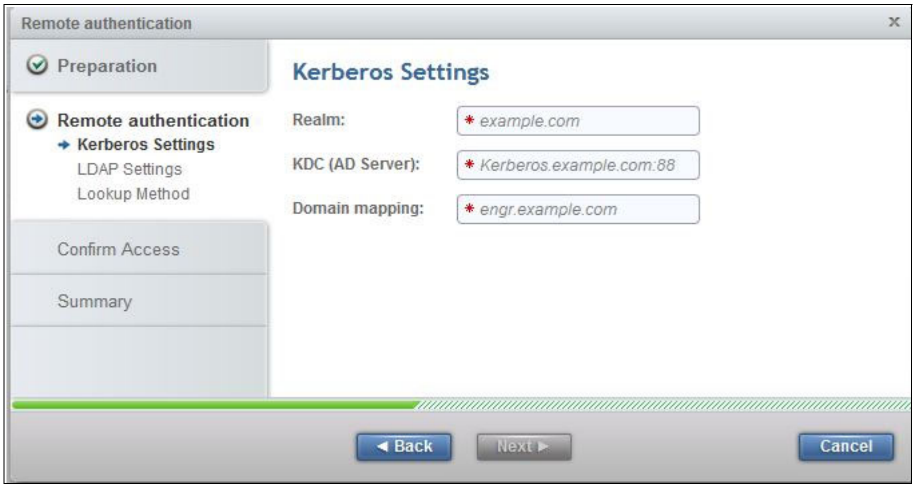

Figure 4-37 Kerberos setting

Complete the following steps in the Kerberos Settings window:

1. Enter the name of the realm setup for this machine.
2. Enter the KDC (AD) server name.
3. Enter the Domain mapping.

Figure 4-38 shows an example of the settings that are used for Kerberos setup. After the settings are chosen, click **Next** to perform a connection check with the KDC (AD server). A green check mark appears if a connection is made. If a connection cannot be made to the server, a red cross appears.

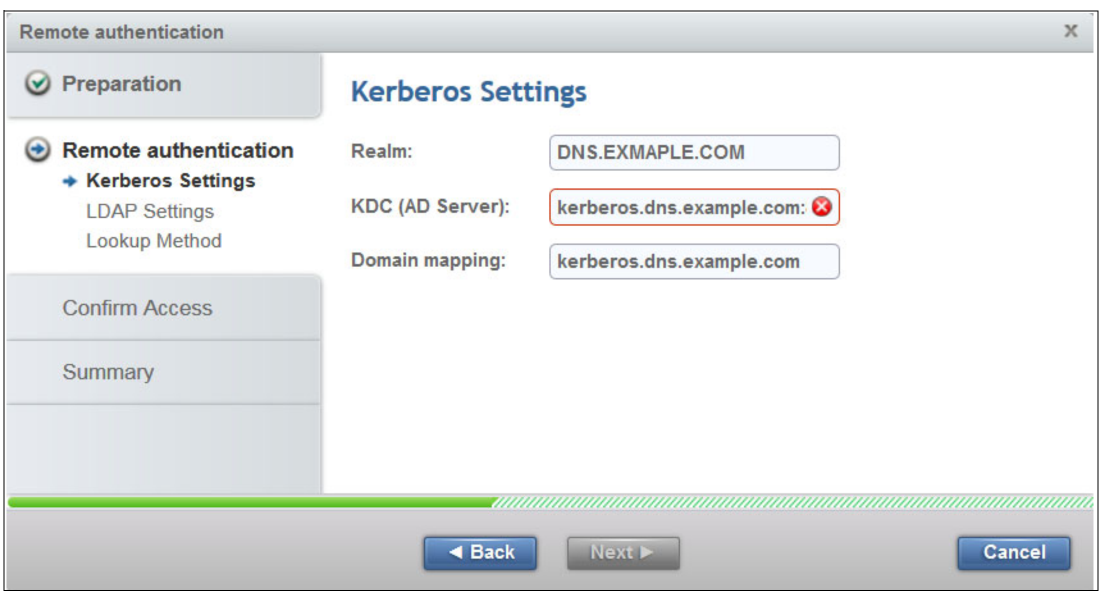

Figure 4-38 Kerberos example settings

LDAP

If LDAP is selected, the LDAP Settings window opens, as shown in Figure 4-39. Complete the required fields in this frame and press **Next**.

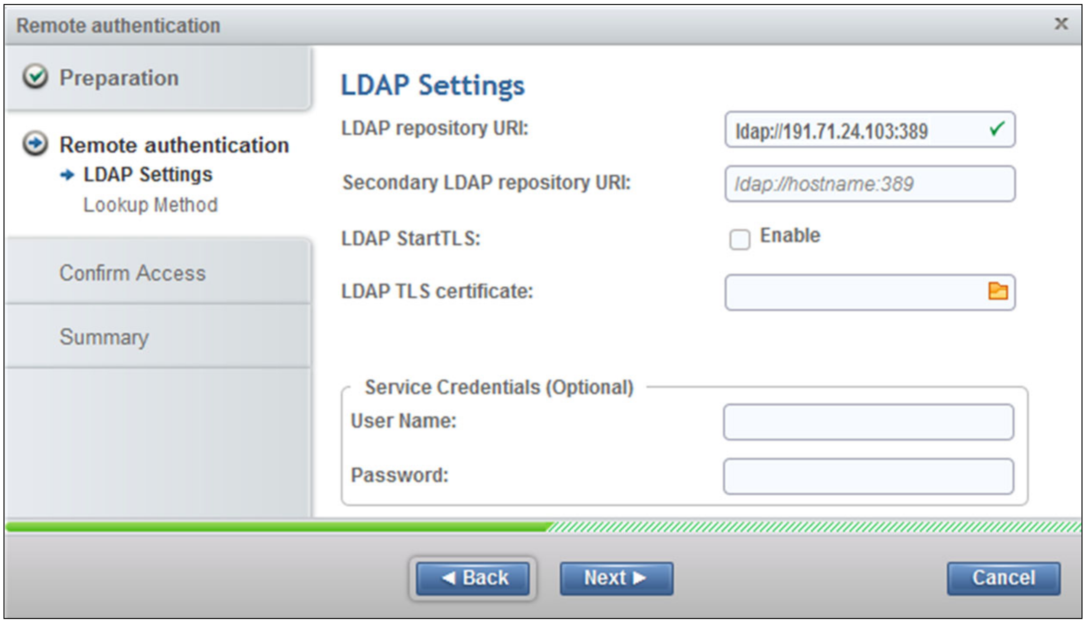

Figure 4-39 LDAP Settings

Configure the following settings on the TS4500 tape library:

- LDAP Repository URI: This field starts with `ldap://` or `ldaps://`and ends with a port number (default for `ldap` is 389 and for `ldaps` is 636).
- Secondary LDAP Repository URI: If you do not have a Secondary Repository URI, leave this field blank.
- LDAP StartTLS: Enabling the LDAP StartTLS starts a normal LDAP session and the TLS (Transport Layer Security) layer (`ldap://` must be used). This selection is optional.
- LDAP TLS certificate: The LDAP TLS certificate is a certificate that the library sends to the LDAP server when the server requires a client certificate. This certificate is required if LDAP StartTLS is enabled. If a TLS certificate is not available, disable LDAP StartTLS.

  When using `ldaps://`, the use of the LDAP TLS certificate is optional.
- Service Credentials: The service credentials include the username and password for the LDAP administrator account. The service credentials are necessary if the LDAP server does not support anonymous access. *Anonymous access* means that any entity can access and view records without being authenticated by the LDAP server.

  The username field must follow LDAP distinguished names (DNs) format to lookup correct service credential user.

  As shown in Figure 4-40 on page 214, the user credentials are

  `CN=Tom,CN=Users,DC=ldapserver1,DC=example,DC=com`

  These generic settings in the example must be updated with the suitable fields and attributes according to your LDAP settings.

Note: Consider the following points:

- If service credentials are not entered, the library uses the user credentials to bind to and query the LDAP server. If this process fails, the library attempts to connect to the LDAP server with an anonymous connection.

  If the TS4500 user does not have sufficient privileges to query the LDAP server, the TS4500 user authentication process fails. When service credentials (with the correct privileges) are supported and entered into the settings, this error does not occur.
- The LDAP connection settings are tested when you select **Next**. A TCP socket is opened to the LDAP-Server. If the settings are correct and communication works, a green check mark appears next to the LDAP repository URI. Then, continue with the next setting. If the LDAP URI is not set up, no connection exists, or if a set up cannot be done, a red cross (+) appears next to the URI.

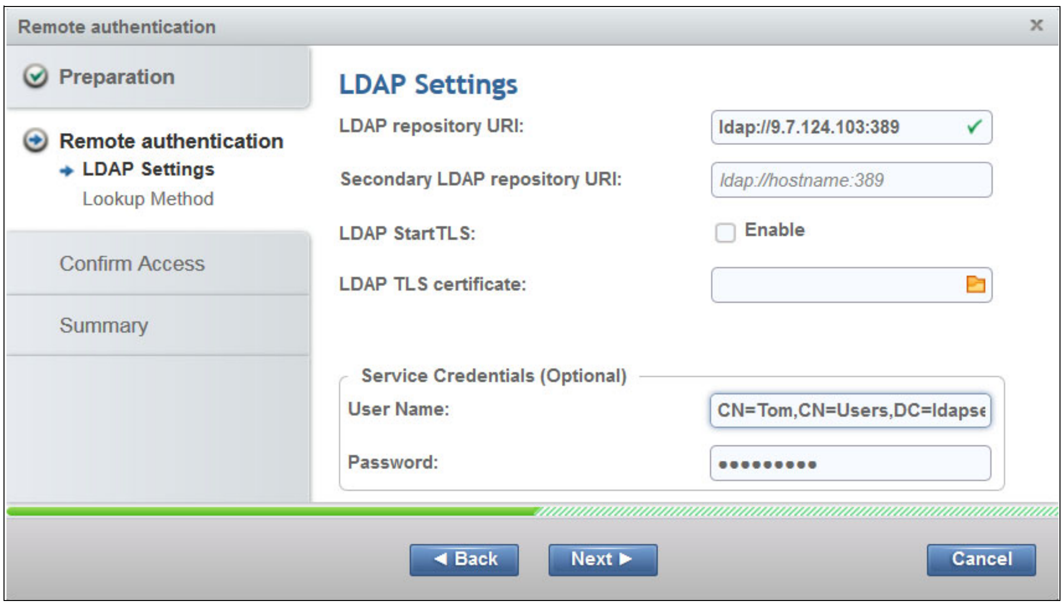

Figure 4-40 LDAP example settings

LDAP lookup methods

Two LDAP lookup methods are available: Simple lookup and Advanced lookup.

Simple lookup

Simple lookup is the default lookup method. It uses group and user LDAP DNs for authentication lookup. After you enter the settings, click **Next** to continue to the following options for the lookup method:

- Group DN specifies the place within the LDAP tree where the groups are located.
- User DN specifies the place within the LDAP tree where the users are located.

Figure 4-41 shows an example of a completed Simple lookup method page that includes an example of the format of the required fields.

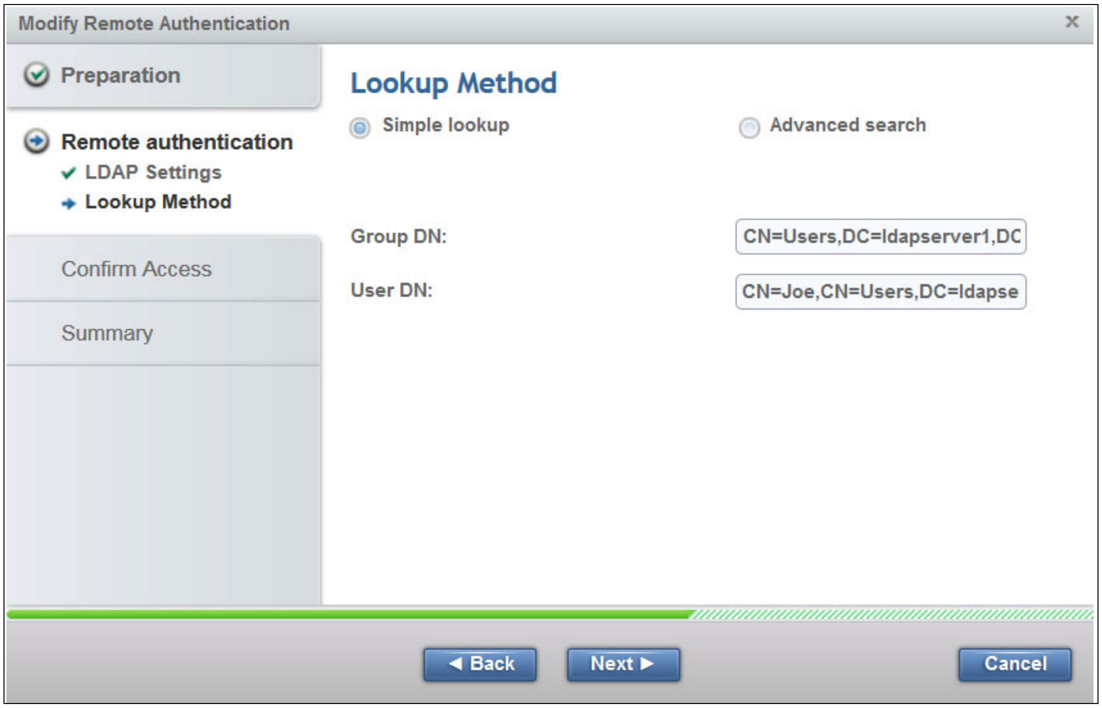

Figure 4-41 Simple lookup method example

Advanced lookup

The Advanced lookup method features more flexible searching and better performance. The Advanced method includes the following search fields, which must be completed:

Base DN: Allows you to customize the Base distinguished name to begin the LDAP search, which begins the search deeper in the LDAP tree for better performance.

Group name Allows you to choose what attribute in the LDAP group accounts is used to associate with a TS4500 role. Default in Active Directory is `cn`.

User name Allows you to customize which attribute in the LDAP user account is used for user names. Default in Active Directory is `cn` or `uid`.

Group Member Allows you to customize the link between the LDAP users and groups. Default in Active Directory is `member` or `memberOf`.

The following fields are optional when selecting Advanced search and are used to improve LDAP search performance:

- User name filter
- Group name filter

Select the advanced search lookup method by selecting the Advanced search option, as shown in Figure 4-42, which also shows how to complete the Advanced search within the LDAP lookup method.

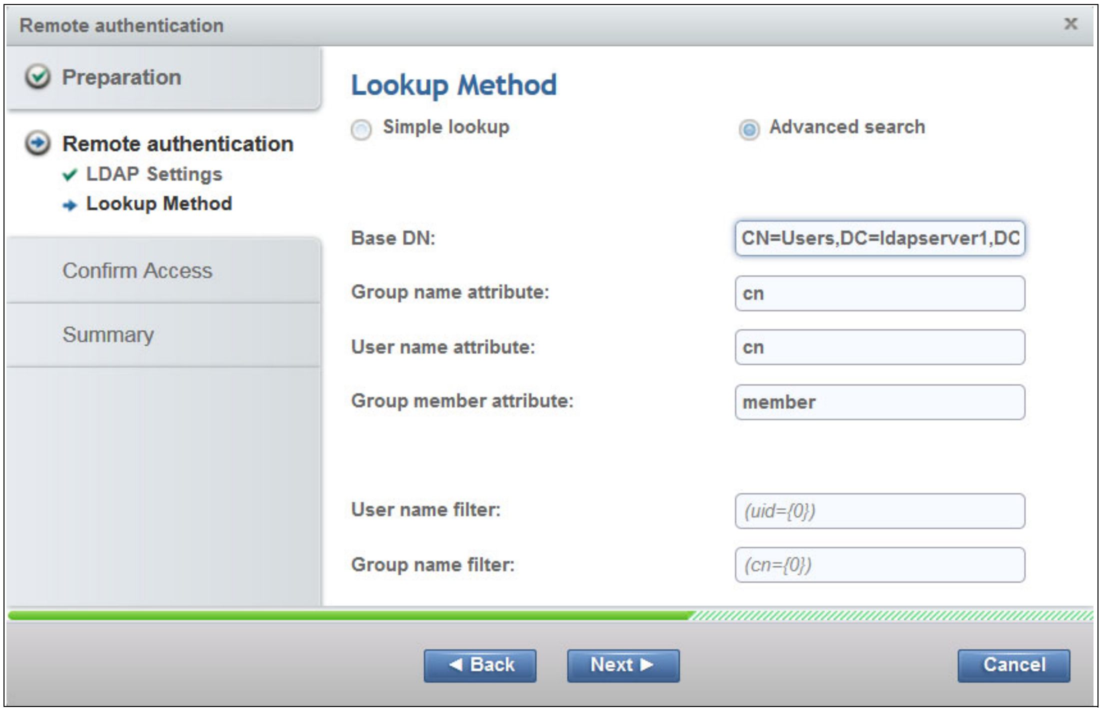

Figure 4-42 Advanced search example

After entering the settings, click **Next** to continue to test the settings are correct and to confirm access to the servers.

In the Confirm Access page, for the Remote UID field, enter a user that was created inside the LDAP server (and the corresponding password). If all settings are correct, remote authentication can be enabled, logs all users out of the management GUI.

If Remote authentication is enabled, local authentication is disabled and no local user can log in.

Resource Access Control Facility

The Resource Access Control Facility (RACF) is the standard security product to manage access control in IBM Z® environment.

The following settings are described for a RACF environment:

- LDAP Repository URI: Must be set to the RACF server IP address and port number (which is the same as LDAP).
- LDAP StartTLS: Not supported
- LDAP TLS certificate: Not supported
- Simple lookup: Not supported
- Advanced search: See the shaded note box that proceeds this list.
- Base DN: `cn=RACF`.
- Group name attribute: `racfgroupid`.

- User name attribute: `racfuserid`.
- Group member attribute: `racfgroupuserids`.

The following settings are optional:

- Service Credentials. If used, the user name must contain the full `racfid`: for example, `racfid=USER1,profiletype=USER,cn=RACF`. For more information, see the proceeding shaded box.
- User name/Group name filters.

> **Note:** These filters can be used with RACF to specify the DN pattern that the tape library must use when determining which groups the user belongs to during the log in process.
>
> When the user/group DN patterns are used in the filter fields, service login credentials are not required. The option to set a DN pattern enables the tape library to support RACF servers, which do *not* allow searching for user information (for example, to get the groups to which the user belongs).
>
> Examples of DN patterns in filter fields:
>
> - User name filter:
>
>   `racfid={0},profiletype=USER,cn=RACF`
>
> - Group name filter:
>
>   `racfid={0},profiletype=GROUP,cn=RACF`

Disabling remote authentication

After successfully enabling remote authentication, clicking **Security** → **Remote Authentication** → **Actions** prompts a Disable remote authentication option. Selecting this option prompts the user to confirm remote authentication disablement. Upon confirmation, log all LDAP users out of the management GUI and revert to local authentication.

Password and Session Policy

Figure 4-43 shows the Password and Session Policy window, which displays the current settings.

Figure 4-43 Password and Session Policy page

If the default password and session policy settings are not strict enough for your company's rules, you can customize the settings on the Password and Session Policy window by selecting **Modify**.

You can modify all of the settings that relate to the session and password policy, as shown in Figure 4-44. These settings can be set back to the default settings by selecting **Reset to Default** in this window.

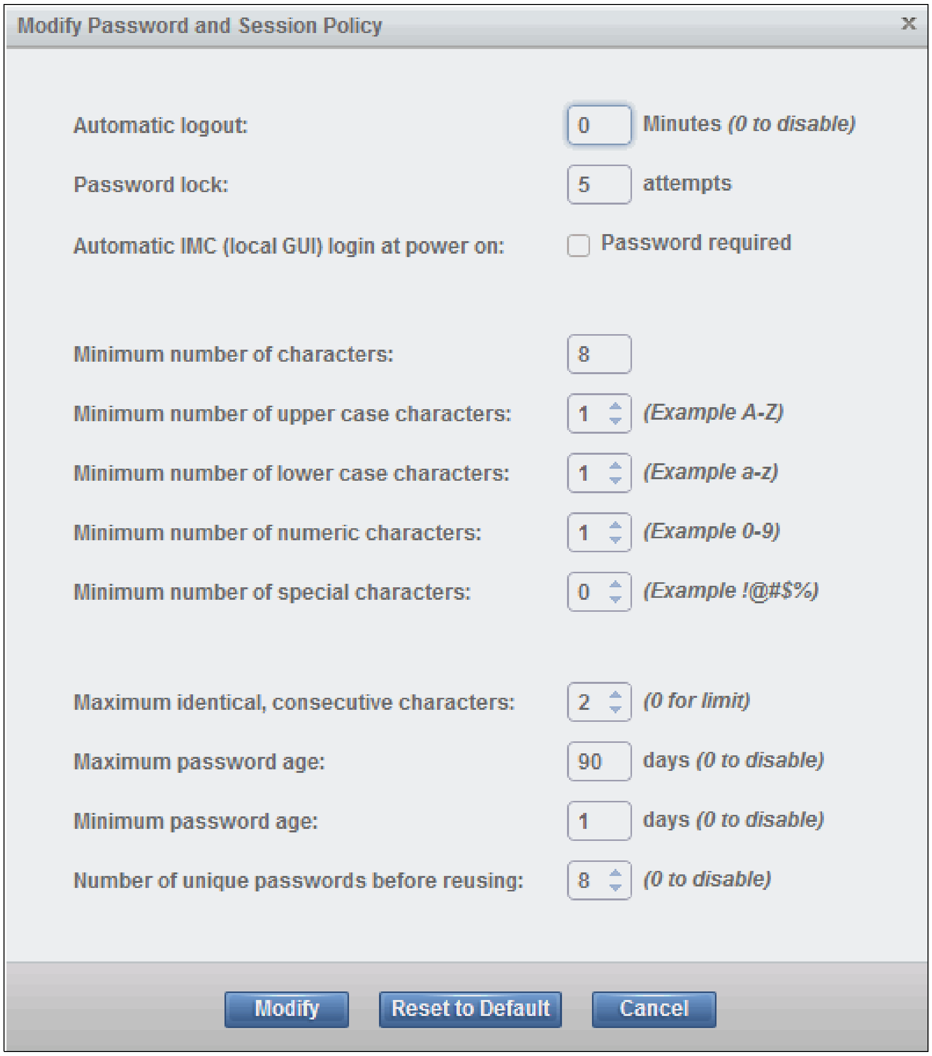

Figure 4-44 Modify Password and Session Policy window

A preset local user that is called `localGUI`, with the role of monitor, is enabled by default with the TS4500 tape library. This default allows a local user to access the integrated management console (IMC) without logging in.

However, if your company policies and procedures prohibit this type of quick access to the System Summary view, you can disable this local user login function. To disable the user login function, select **Settings** → **Security** → **Password and Session Policy**, and set the Automatic IMC (local GUI) login at power on setting so that a password is required.

Secure communications

Use this option to manage encryption settings for the GUI and command-line interface (CLI) communications with the library.

The TS4500 tape library is secured with a Secure Sockets Layer (SSL). SSL is a protocol for encrypted (secure) transmission through the internet.

Use the Secure Communications page to configure the SSL settings (see Figure 4-45).

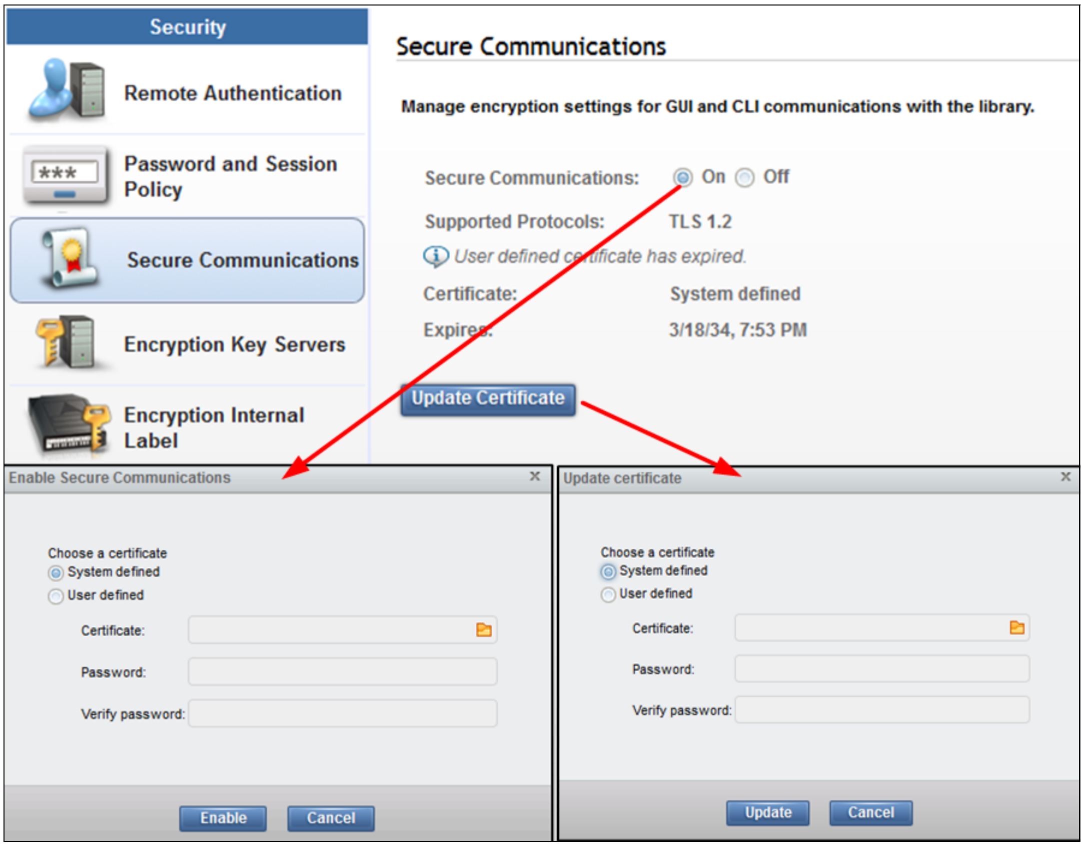

Figure 4-45 HTTPS

Use the Secure Communications window to enable or disable SSL or to update a certificate. Selecting **Enable** or **Update** provides the option to upload an SSL web certificate. This window also displays the type of certificate that is used and the upload date and expiration date.

SSL is a cryptographic security system that uses the following keys to encrypt data:

- A public key that is known to everyone
- A private key that is known only to the recipient of the message

Many websites use this protocol to obtain confidential user information, such as credit card numbers. By convention, URLs that require an SSL connection start with `https` instead of `http`. HTTPS stands for Hypertext Transfer Protocol Secure.

Secure communications on

When the secure communications function is on, the data that is exchanged between the TS4500 tape library and the browser or CLI is encrypted. In the browser, you must add and confirm security exception to login. For CLI, you must use the `--ssl` option to run CLI commands.

Secure communications off

When the secure communications function is off, your browser or CLI and the TS4500 tape library communicate with unencrypted data transmissions. Your browser uses a TLS certificate to verify the validity of the TS4500 tape library.

Updating web certificate

SSL web certificates are small data files that digitally bind a cryptographic key to an organization's details. When an SSL web certificate is installed on a web server, it activates the padlock and the `https` protocol (over port `443`) and allows secure connections from a web server to a browser. This option can be used to update a certificate or replace an expired certificate.

The following options are available when you enable secure communications or update a certificate:

- System-defined

  The SSL cipher specification list system value is read-only. Its values are automatically modified to contain the list of cipher suites that are supported by the system SSL. If you use this option, the SSL cipher specification list system value is automatically updated with new cipher suite capabilities when you install or upgrade to a future release of the firmware system.

- User-defined

  If a certificate exists, its identifying information is displayed in the Certificate field. You can use this certificate, or click the folder icon to upload another certificate.

  > **Note:** If you modify a certificate that was uploaded, you must re-select the PEM file by clicking the folder icon and specifying the file, even if the name is still displayed in the Certificate field. If the certificate is encrypted, you must re-enter the password.

Consider the following points regarding certificates:

- RSA keys are recommended. No specific key size is required.
- Certificates must use AES encryption. DES is not supported.
- Certificates must be in PEM format. They can be self-signed or CA-signed, but must contain the certificate and the private key. If the private key is encrypted, you must enter the password for the private key.
- Certificates with the SHA256 signature algorithm are supported. Certificates with SHA1 or MD5 hash signatures are not allowed.
- Both wildcard and multi-domain (SAN) certificates are supported. A wildcard certificate allows unlimited subdomains to be protected with a single certificate. A SAN certificate allows for multiple domain names to be protected with a single certificate.

If you are using CA-signed certificates on a TS4500 tape library configured with multiple IPs, it is strongly recommended to include each IP used for web GUI access in the certificate. If an IP address or DNS name is not specified in the certificate, the message, “Your connection is not secure”, is displayed after the web server restarts. Click **Add Exception** → **Confirm Security Exception** to use that certificate.

For more information about SSL, see IBM Documentation.

Encryption key servers

Use the Encryption Key Servers page to manage the key servers that use an encryption key.

Since Firmware 1.8.x, secure communication between the library and key server is supported. The needed certificate must be in PEM-Format. For more information, see “Add Encryption Key Server on the TS4500” on page 180.

From the Add Server tab, add each encryption key server to use on the logical libraries configuration of this TS4500, as shown in Figure 4-46. This action makes the IP address of the server available when you set up library-managed encryption on the logical library.

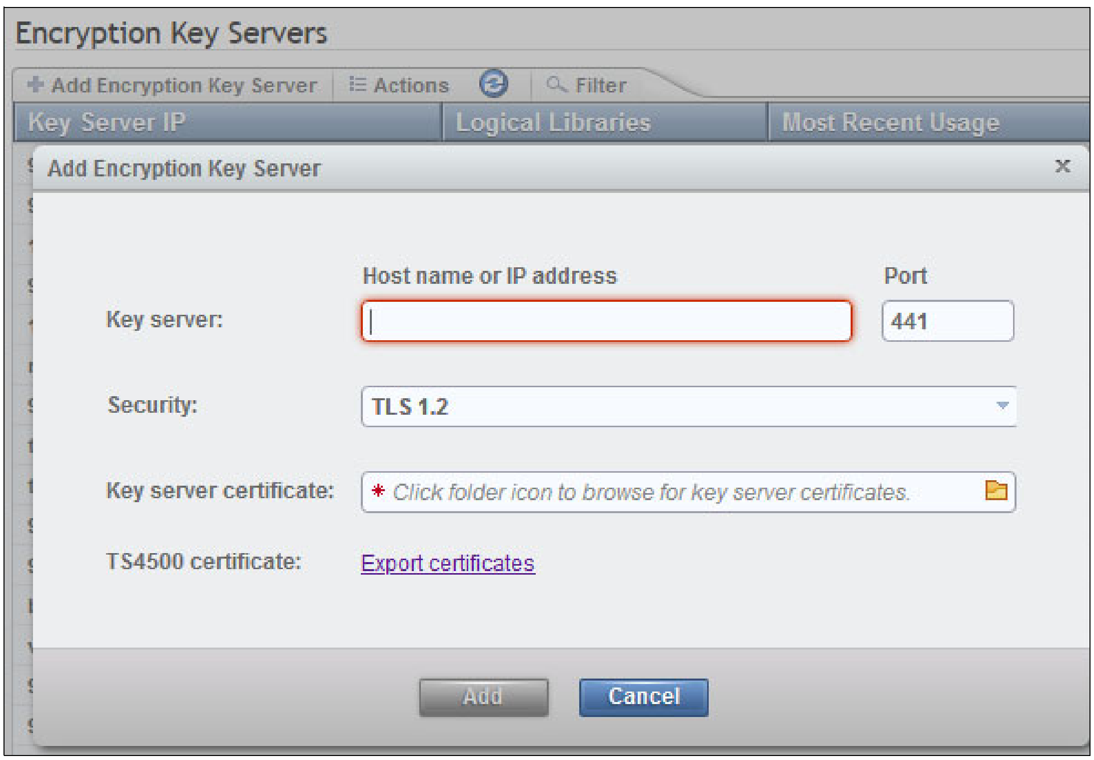

Figure 4-46 Encryption Key Servers page

Encryption is managed at the logical library level. All encryption-enabled drives that are assigned to a logical library use the same method of encryption. For more information about methods of encryption, see 4.7.2, “Creating Logical Library window” on page 284 and Chapter 3, “Encryption” on page 173.

The Add Encryption Key Server window requires you to first check with an Internet Control Message Protocol (ICMP) ping to ensure that the server is available when you add the server. Click **Ping** (see Figure 4-46) when you add a server to the list. Encryption key servers can also be added when you modify a logical library.

Note: The TS4500 supports the configuration of four encryption key servers for each logical library.

Run Diagnostics option

If you want to verify the functionality of all of the network-attached devices, you can run diagnostics on the devices. Select **Actions** → **Run Diagnostics**. Then, select the drive or select drive to test and select **Run Diagnostics**.

The diagnostics process runs for a few minutes. When the diagnostics complete, the window shows a check mark for devices that passed the test, as shown in Figure 4-47, or a red error icon for devices that failed the test.

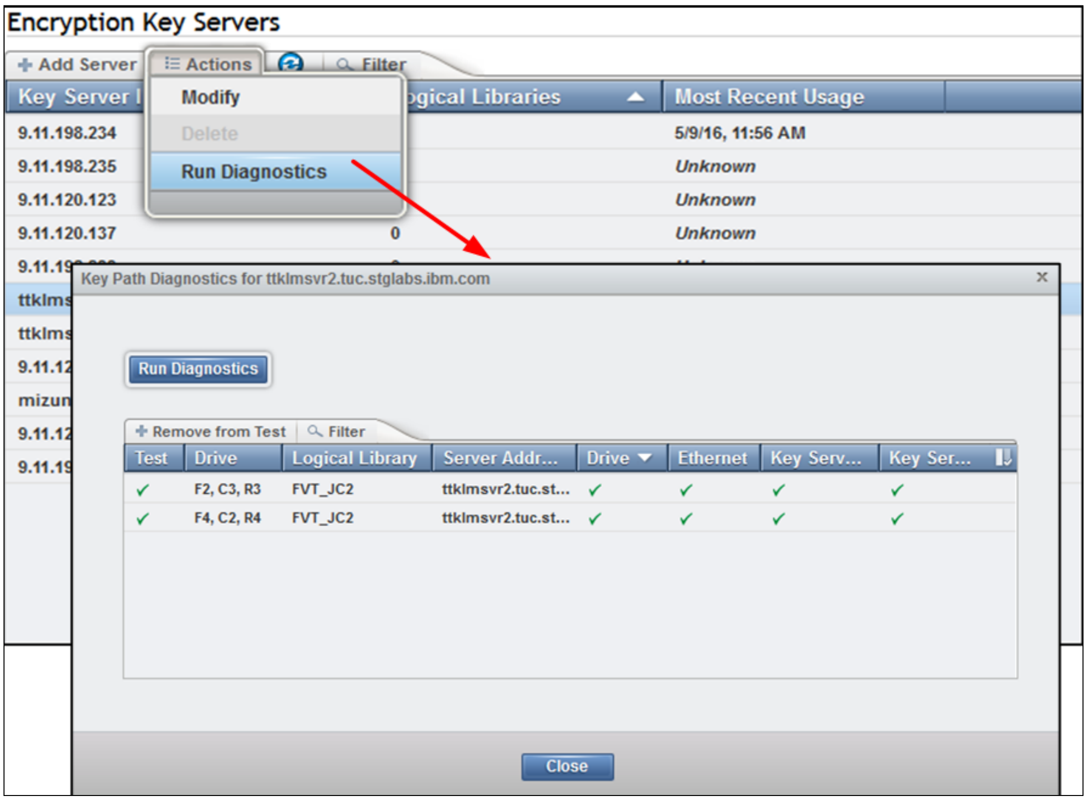

Figure 4-47 Encryption Key Servers diagnostics

Note: When you use the Run Diagnostics function, a warning icon might appear in the Ethernet column if you disabled Internet Control Message Protocol (ICMP) requests.

At least one key server IP address must be selected for the Run Diagnostics option to display on the Actions menu (or the right-click drop-down menu). If no key server IP address is added (or selected), the Run Diagnostics option does not display on the Actions menu. This server must be configured on a logical library for diagnostics to run.

ICMP requests

Internet Control Message Protocol (ICMP) is a messaging protocol (external to the TS4500 tape library) that sends error messages that a requested service is not available or that one of the network-attached devices cannot be reached. ICMP requests are disabled or enabled on the server on which the IBM Security Guardium Key Lifecycle Manager (GKLM) runs.

The disabled status of ICMP requests can trigger a warning in the Ethernet column when the Run Diagnostics function starts.

If ICMP requests are disabled, the server that the Encryption Key Manager runs on does not return a ping when the diagnostics are run, which triggers a warning in the Ethernet column of the Run Diagnostics table. If a warning appears in the Ethernet column, but no warning appears in the Key Server Path column, ignore the warning in the Ethernet column. If warnings appear in the Ethernet column and the Key Server Path column, the failure might be on the machine, or the network connections between the library and that machine.

Encryption Internal Label

Use the Encryption Internal Label page to create, change, or delete mappings from the cartridge key labels to the key-encrypting labels.

When the Encryption Internal Label option is configured, the encryption-enabled tape drive automatically derives the encryption policy and key information from the metadata that is written on the tape volume by the TS4500 tape library. Mapping the cartridge key labels to the key-encrypting labels enables the TS4500 tape library to apply the same encryption policy for both types of labels. The Encryption Internal Label is a way for the system to share encryption policies.

For a TS4500 with only a single drive type that is installed, either Linear Tape-Open (LTO) or 3592, only one option is shown on the Security page, as shown in Figure 4-48.

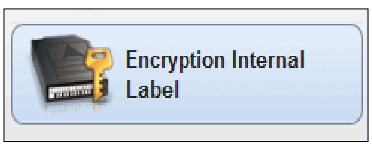

Figure 4-48 Single drive type Encryption Internal Label

Encryption is managed at the logical library level. All encryption-enabled drives that are assigned to a logical library use the same method of encryption. For more information about methods of encryption, see 4.7.2, “Creating Logical Library window” on page 284 and Chapter 3, “Encryption” on page 173.

The examples that are shown in the next sections show the display for a mixed drive type library.

Encryption Internal Label 3592

To configure key label mapping, select **Encryption Internal Label 3592** and then, select the **Create Mapping** tab, as shown in Figure 4-49. Select the method to use and then, select the key labels, as required.

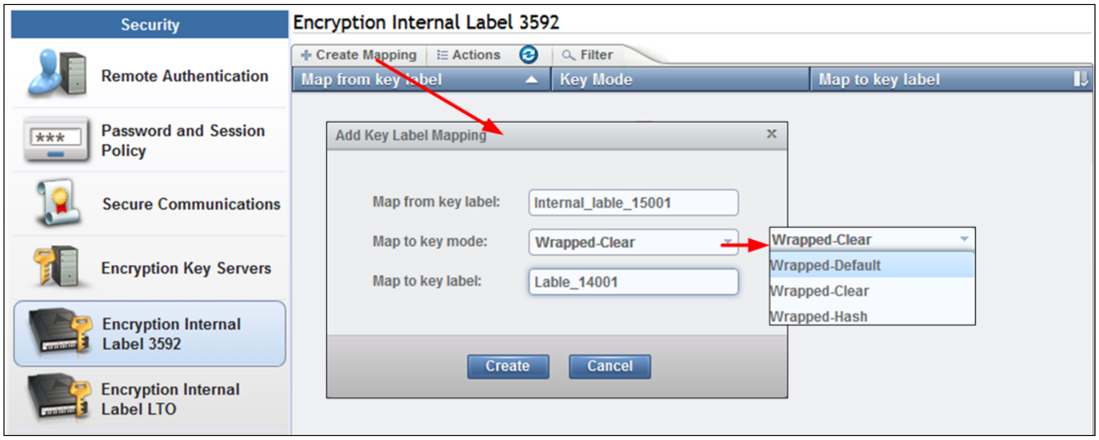

Figure 4-49 Add or modify key label mapping for 3592

The following values are available for the “Map to key mode” field:

- Wrapped-Default: The map to key encryption method is configured by using the Encryption Key Manager default. (This option is for 3592 cartridges only.)
- Wrapped-Clear: The externally encoded data key (EEDK) is referenced by the specified key label. (This option is for 3592 cartridges only.) The Wrapped-Clear method is typically specified when encrypted volumes are kept in-house where each keystore references the keys by using the same key labels.
- Wrapped-Hash: The EEDK is referenced by a computer value that corresponds to the public key that is referenced by the specified key label. (This option is for 3592 cartridges only.) The Wrapped-Hash method facilitates exchange with a business partner or when volumes are sent to a disaster recovery site where the key labels might differ for the same key.

Encryption Internal Label LTO

To configure key label mapping, select **Encryption Internal Label LTO** and then, select the **Create Mapping** tab, as shown in Figure 4-50. Select the method to use and the key labels, as required.

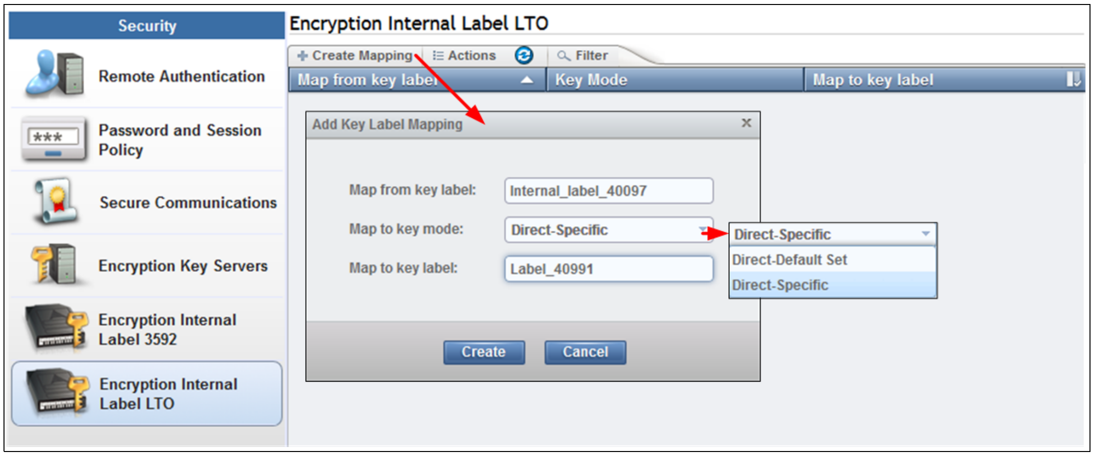

Figure 4-50 Add or modify key label mapping for LTO

The following values are available for the “Map to key mode” field:

- Direct-Default Set: The map to label is determined from the encryption key manager. The label was configured at the encryption key manager, and the key label field is left blank. (This option is for LTO cartridges only.)
- Direct-Specific: The specified key label references a symmetric data key. (This option is for LTO cartridges only.)

4.3.5 GUI Preferences

Select GUI Preferences on the Settings menu to access the GUI Preferences page to control the behavior of the navigation dock. Use the Navigation page in the TS4500 Library to enable or disable the animation of the navigation dock in the GUI. Figure 4-51 shows example of animation enabled.

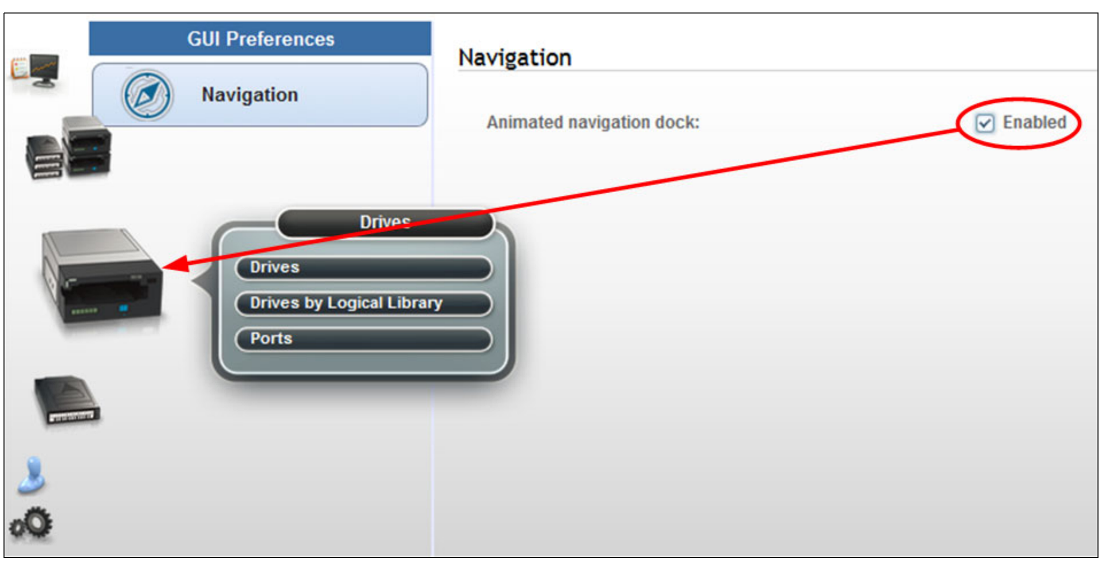

Figure 4-51 Enabled navigation animation

You can turn off the animation if the enlarged icons obscure any information in the window. Also, turning off the navigation animation helps increase performance when you are remotely connecting to the TS4500 over a slow network connection. Figure 4-52 shows example of animation disabled.

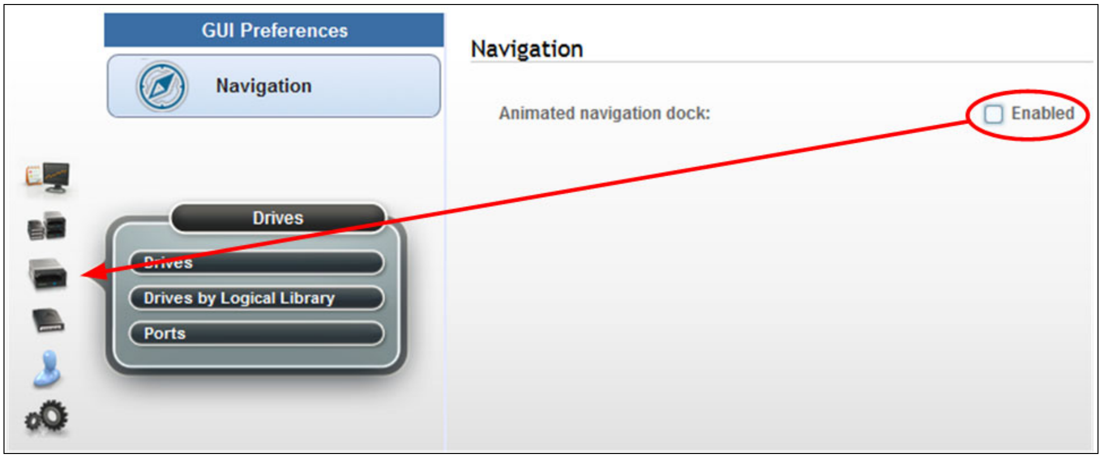

Figure 4-52 Disabled navigation animation

4.4 Access menu

Use the TS4500 management GUI pages (which is under the Access icon), as shown in Figure 4-53, to view, create, and assign users and their roles.

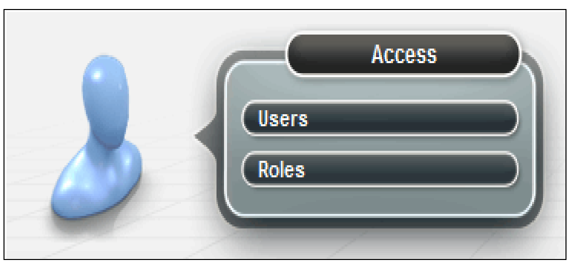

Figure 4-53 Access menu

Note: The options and actions under the Access icon are available only to users with the administrator role.

4.4.1 Users

Administrators can create and manage users, map users to a role, and view which users have active sessions and how many connections they have from the Users page.

Note: With the TS4500 tape library, password protection is always enabled and all users are required to sign in with a user name and password.

Selecting **Users** from the Access menu gives an overview of all configured users, as shown in Figure 4-54.

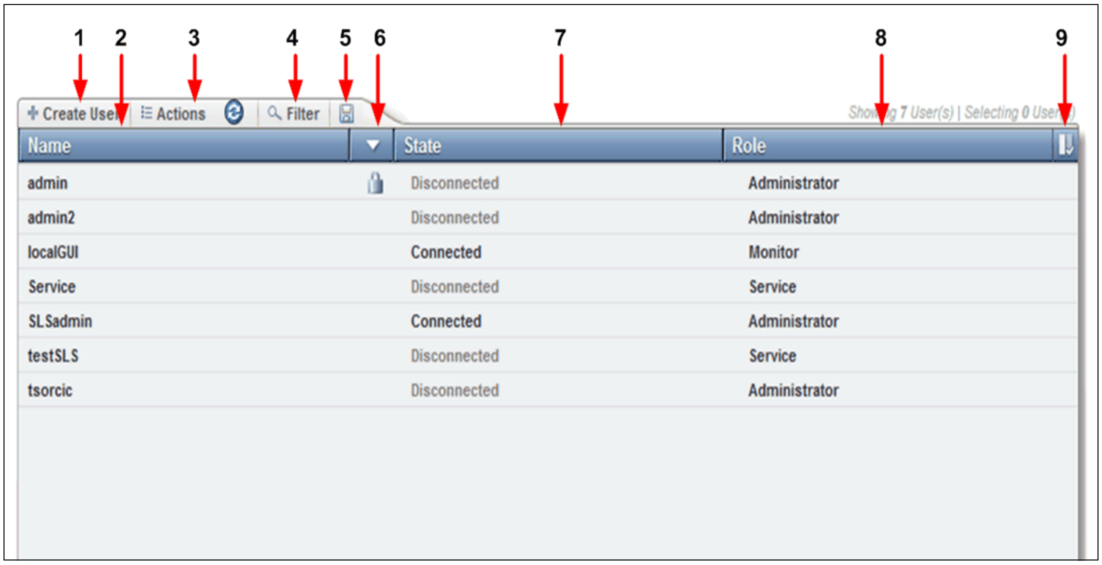

Figure 4-54 Create User tab

Table 4-4 lists the fields that are available on the Create User tab. The numbers in the table correspond to the numbers that are shown in Figure 4-54.

Table 4-4 User window

Number	Description
1	Create User tab to add a user
2	Name of user
3	Actions tab
4	Filter user or search tab
5	Save user list to file
6	Locked state of user
7	Connected state of user
8	User role
9	View menu options

The actions that are available from the Users page differ, depending on whether the library is configured for local or remote authentication. Table 4-5 lists the available actions.

Table 4-5 Local and remote authentication actions

Actions	Local authentication	Remote authentication
Create User	X	
Map To Role	X	
Modify Email	X	
Reset Password	X	
Delete User	X	
Disconnect	X	X
Connections	X	X
Properties	X	X

Local authentication

With local authentication, each TS4500 tape library maintains a separate database of user names with corresponding passwords and roles. Therefore, user management must be performed on each library.

When local authentication is enabled, the Users page shows all users and their state (connected or disconnected). The number of active connections is displayed next to the user state. It is also possible to view and modify email addresses when local authentication is enabled.

Remote authentication (LDAP)

With remote authentication, security tasks are centralized and user management can be performed from a single interface, without logging in to the TS4500 tape library. When remote authentication is enabled, authentication requests are passed to a remote authentication server (LDAP) that verifies the user's name and password.

On the LDAP-Server, this user is associated with a number or groups. If one of these groups matches a customized role within the TS4500, the user is granted the permissions that are associated with the role on the library.

Important: The name of the group on the remote authentication server must be the *same* as the name of the role on the TS4500 tape library for the group to be matched to that role.

We support characters [a-z][A-Z][0-9] and "-" and "_" with a maximum length of 50 characters. Roles with spaces are *not* supported.

The Create option is disabled on the Users page when remote authentication is enabled because you must create the user on the remote authentication server. In addition, modifying a user's email, password, and user group must be performed on the remote authentication server. Only connected users are shown on the User page, as shown in Figure 4-54 on page 228.

For more information about enabling remote authentication, see "Enable remote authentication" on page 210.

IMC access by a local user

In some environments, it is possible to rely on the physical security of the data center as the default operating mode. For this scenario, the TS4500 tape library allows a local user to access the Integrated Management Console (IMC) without logging in, enabling quick access to the System Summary view of the TS4500 management GUI.

A preset local user called `localGUI`, with the role of monitor, is enabled by default with the TS4500 tape library. This is the only local user that is valid when remote authentication is enabled.

If the localGUI user is deleted, all users are required to log in at the IMC.

To disable this local user login function and show the login window at the IMC, click **Settings** → **Security** and then, select **Access Rules**.

Creating users

You can create up to 80 users, each with a unique user name and password.

Each user account is mapped to a role that defines the pages that the user can view and the actions that the user can perform. Each user can be mapped to only one role, but multiple users can be mapped to any role.

For more information about the user roles and how to configure them, see 4.4.2, "Roles" on page 234.

To create a user, select **Create User** in the Users window. The Create User window opens, as shown in Figure 4-55. Enter the user name and role. The user receives the role that is chosen from the Role drop-down list. If you do not want the user to receive a default role, ensure that the role is first set up so that it can appear on the drop-down list.

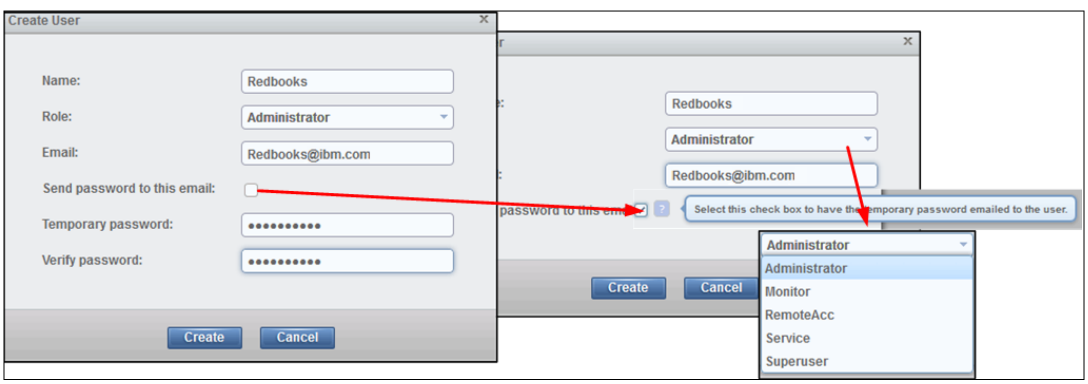

Figure 4-55 Create User option

For more information about creating and managing roles, see 4.4.2, "Roles" on page 234.

A user can receive the temporary password by using one of the following methods:

- If the Send password to this email option is selected, you must specify a valid email address for a user so that the user can receive passwords and email notifications. You also must set up email notifications first (for more information, see "Email Server" on page 205).

  When a user is created or when a user's password is reset, the temporary password is automatically sent to the email address of the user, rather than to an administrator. A user with a valid email address can also be selected as a recipient of email alerts from the Notifications page.

- If an email address is not specified, an administrator is responsible for generating a new password or resetting a password and informing the user.

The password that you set must conform to the rules that are defined on the Password and Session Policy window. For more information, see “Password and Session Policy” on page 218.

> **Important:** Because a user’s password expires after the user is created, a new user is prompted to reset the password the first time that the new user logs in.

After all fields are completed, select **Modify** to receive a message that confirms that the user is created and a reminder that the password is only temporary (see Figure 4-56).

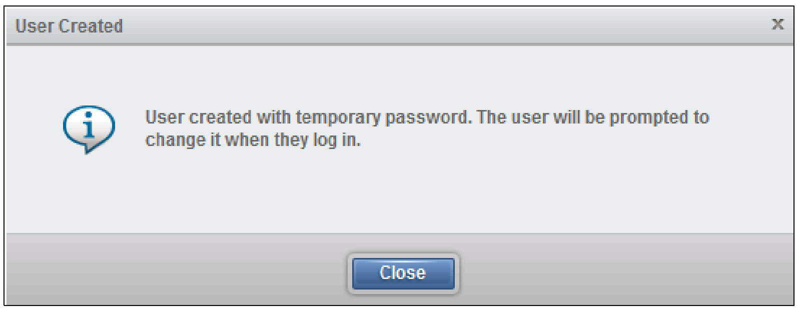

Figure 4-56 User Created message

> **Note:** It is highly advised to create at least two separate Administrator users because only the Administrator can reset or unlock a user password.

First-time login

When a user logs in for the first time, the user is presented with the login window. The user must enter their username and temporary password to start the login process, as shown in Figure 4-57.

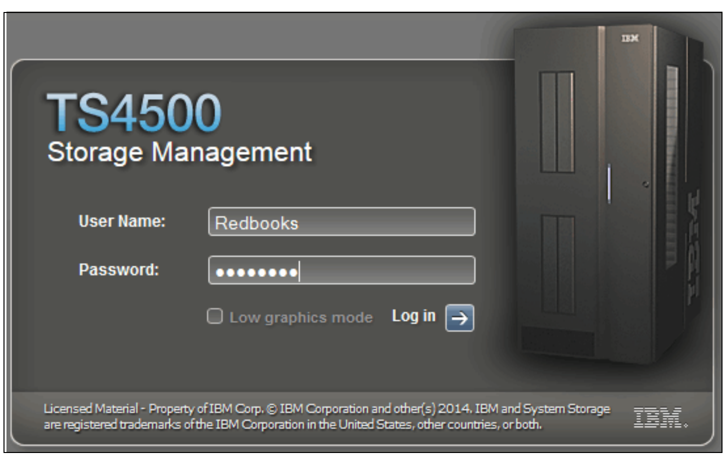

Figure 4-57 User login

The user is then presented with the change password window, as shown in Figure 4-58.

Figure 4-58 Change password

The new password must conform to rules defined in the Password and Session Policy. For more information, see "Password and Session Policy" on page 218.

Note: The username and password are case-sensitive.

After the user sets their password, the user logs in with the new password in the initial login window, as shown in Figure 4-57. After the user logs in with the new password, the user can access the TS4500 management GUI with access to the assigned role.

Modifying a user

An administrator can perform several actions on a user by selecting the user and using the Actions menu, as shown in Figure 4-59.

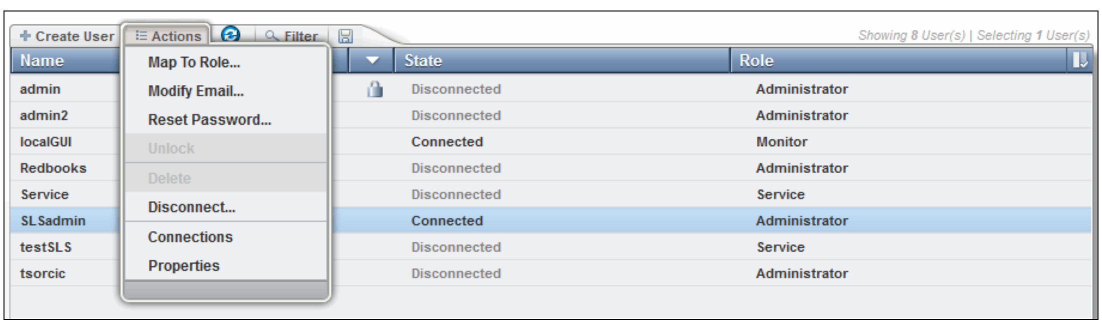

Figure 4-59 Actions

The following functions are available on the Actions menu:

- Map to Role: The administrator can map a new role to a user. If the user is logged in, this action logs out the user.
- Modify Email: The administrator can modify or add an email address to a user.
- Reset Password: The administrator can reset a user password and provide the user with a temporary password. The user must change this password to a new password by using the same procedure that is described in "First-time login" on page 231.

- Unlock: If a user is locked out because the user exceeded the maximum number of unsuccessful login attempts, an administrator can unlock the user. The locked user shows a locked icon on the user list, as shown in Figure 4-60.

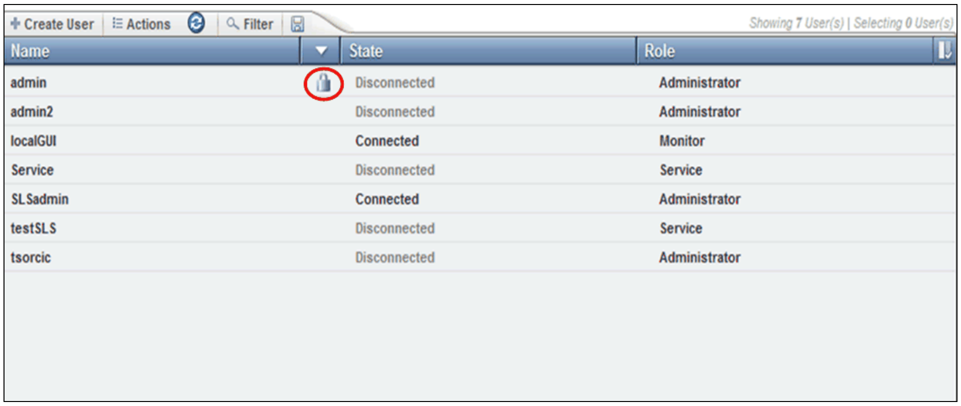

Figure 4-60 Locked user

If the user lost their password, any administrator can reset the password. If the user is locked, the administrator must unlock the user first.

- Delete: An administrator can delete a user. The user must be in a disconnected state to be deleted.
- Disconnect: An administrator can disconnect a user from the TS4500.
- Connections: This option lists the IP addresses that were used by the selected user to log in to the system. This option shows whether a user is being used multiple times and from what address.
- Properties: This option displays the properties of the user, the state of the user, and the user's last connection, as shown in Figure 4-61.

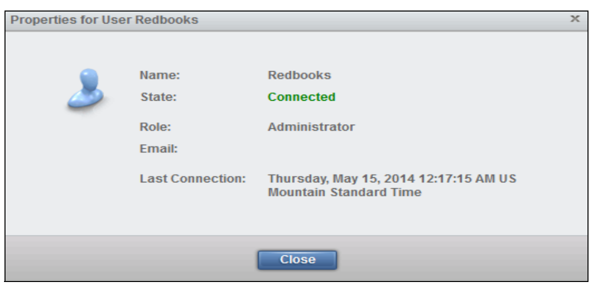

Figure 4-61 User properties

Locked out

If you are locked out of the TS4500 management GUI, you can revert to the default administrator user name and password by using the reset button located on the panel near the pause button of the TS4500 tape library. This button resets the user name and password to the default logon settings for 15 minutes.

Use this reset button only for emergencies; for example, if the remote authentication server is down or if the administrator loses or forgets their password and no administrators can access the system. For more information about this process, see "Access Recovery" on page 55.

4.4.2 Roles

Roles define the pages that users can view and the actions that they can perform. Library administrators manage roles and the users that are assigned to them. Figure 4-62 shows the Role window.

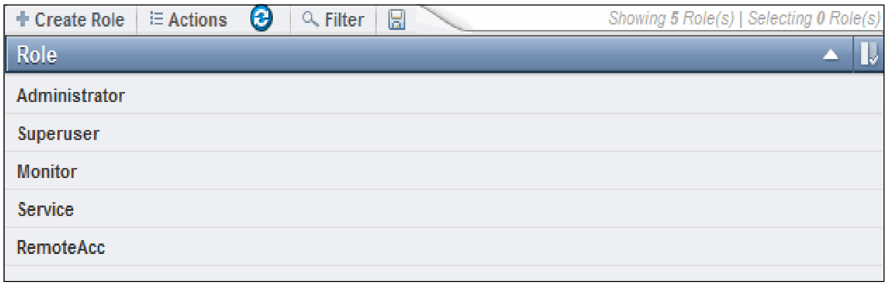

Figure 4-62 Role window

Preset roles

The TS4500 tape library has four preset roles, each with a different level of access to the TS4500 management GUI. It is also possible to create up to 16 roles, each with a custom name and one of the preset permission levels. All roles have access to all logical libraries.

You cannot rename, delete, or modify the access permissions for any of the preset roles.

The following roles are preset default roles:

- Monitor: Users that are mapped to this role can view all physical and library data, but they cannot view user accounts or security settings. This role is useful for library operators.
- Superuser: Users that are mapped to this role can view all pages and perform library tasks, but they cannot manage users, modify security settings, or access service-related functions.
- Administrator: Users that are mapped to this role can perform all library tasks, including managing access and security, but they cannot access service-related functions.
- Service: Users that are mapped to this role can view all pages that are available to a monitor role. Additionally, this role can perform service-related functions, such as updating firmware, downloading logs, calibrating library components, and performing diagnostic tests. This role is useful for IBM service support representatives (SSRs).

  When this user logs in to the TS4500 management GUI, certain tables display additional information, and several of the available actions differ from the actions that are available for the other roles.

Custom roles

Users that are mapped to a role with a customized name can view all of the pages that are available to the selected preset permission level of the custom role.

> **Note:** The permissions of a custom role with service-level access cannot be modified, even from the TS4500 CLI.

Creating or deleting a role

Complete the following steps to name new roles and assign their permission by selecting from the preset role permission levels:

1. Create roles by selecting **Access** → **Roles**.
2. Select the **Create Role** tab.
3. Enter a name for the role.
4. Select one of the preset role permission levels from the Duplicate permission from list box (see Figure 4-63). All roles can access all logical libraries in the TS4500 tape library.

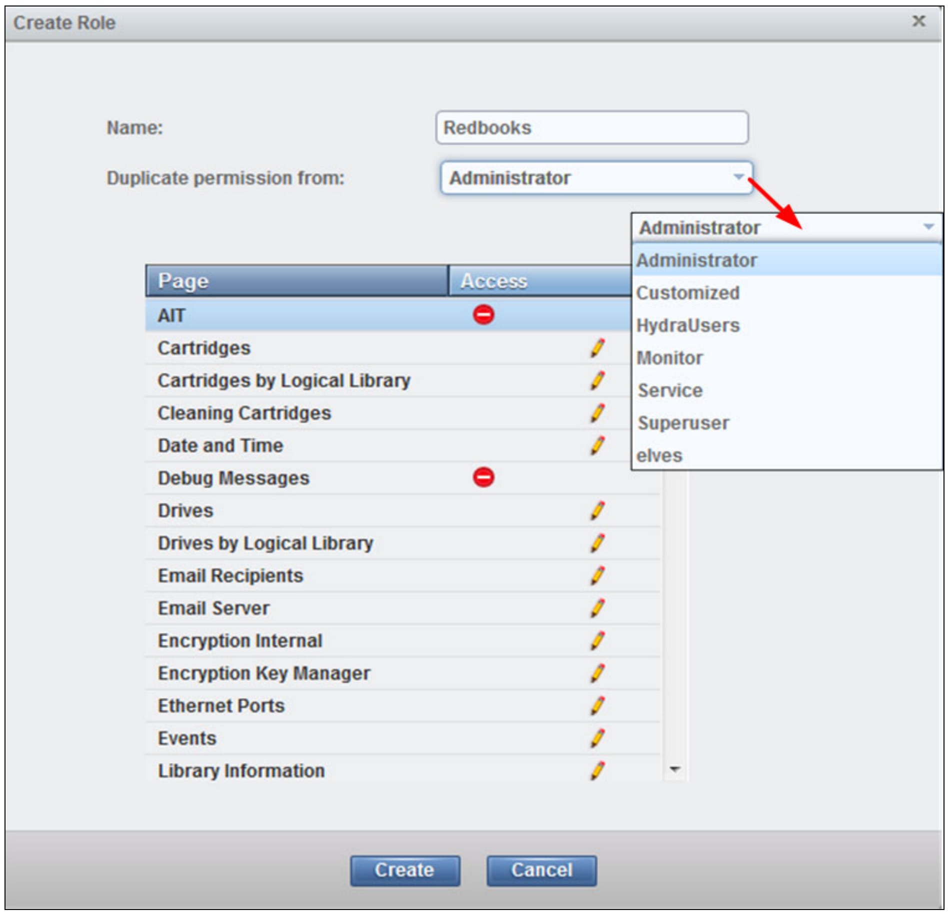

Figure 4-63 Creating a role

To remove a role, select the role. Then, click **Access** → **Roles** → **Actions** → **Delete**. You can delete a role only if no users are assigned to the role.

You can also use the TS4500 CLI to create, modify, and delete roles, and to customize the permissions of a role.

Note: You must create a custom role before a user can be assigned to it.

Modifying role permissions

You can modify customized role tasks from the TS4500 CLI only. However, the role must be created first by using the process that is described in “Creating or deleting a role” on page 235.

Ensure that the TS4500 CLI is installed. Open a command prompt in the same directory as the `.jar` file.

The CLI command features the following format:

```
java -jar TS4500_CLI.jar "C:\[pathname_of_ts4500_cli.jar]" -ip [LCC ip] -u
[username] -p [password] --[cli_command]
```

This command is useful to gather current role permissions that you can modify.

> **Note:** If SSL is enabled on the TS4500 tape library, you must add the **`--ssl`** parameter to the command after the password, as shown in the following example:
>
> ```
> java -jar TS4500CLI.jar -ip [LCC ip] -u [username] -p [password] --ssl
> --[CLI_command]
> ```

viewRolePermissions

The **`viewRolePermissions`** command displays a list of all permissions for a specified role. For example, we set up a role that is called `Redbooks`, which we want to save to a file that is called `role.txt`, as shown in Example 4-1.

Example 4-1 viewRolePermissions command

```
C:\TS4500>java -jar TS4500CLI.jar -ip 10.1.121.99 -u Redbooks -p RedBOOks --view
RolePermissions Redbooks > role.txt
                           Action,   Access Level
                              AIT,      No Access
                       Cartridges,         Modify
      Cartridges by Logical Library,         Modify
               Cleaning Cartridges,        Modify
                    Debug Messages,      No Access
                            Drives,         Modify
          Drives by Logical Library,         Modify
               Email Notifications,         Modify
                  Email Recipients,         Modify
               Encryption Internal,         Modify
            Encryption Key Manager,         Modify
                    Ethernet Ports,         Modify
                            Events,         Modify
               Library Information,         Modify
                Licensed Functions,         Modify
                 Logical Libraries,         Modify
           Management GUI Behavior,         Modify
                    Master Console,      No Access
                    Password Rules,         Modify
                  Position Control,      No Access
             Remote Authentication,         Modify
                             Roles,         Modify
                 SNMP Destinations,         Modify
                SNMP Notifications,         Modify
                        Scan Speed,      No Access
```

```
                Secure Socket Layer,         Modify
                       Service Port,      No Access
             Syslogs Notifications,         Modify
                             System,         Modify
              System Date and Time,         Modify
                              Tasks,         Modify
                              Users,         Modify
                      VOLSER Ranges,         Modify
  VOLSER Ranges by Logical Library,         Modify
```

The output of this command is saved to our file, `role.txt`. You can now modify this file and change roles to suit your requirements. The output from this command can also be copied to another text file and modified.

setRolePermissions

The **setRolePermissions** command sets up custom roles to the permissions for a specified role. First, edit the text file that was created by using the **viewRolePermissions** command. Then, set the new permissions by using the **setRolePermissions** command, as shown in Example 4-2.

Example 4-2 setRolePermissions

```
C:\TS4500>java -jar TS4500CLI.jar -ip 10.1.121.99 -u Redbooks -p RedB00ks --set
RolePermissions rolenew.txt -role Redbooks
The permissions were updated successfully
```

To confirm that the role is changed, run the **viewRolePermissions** command, as shown in Example 4-3.

Example 4-3 Confirming that role is changed

```
C:\TS4500>java -jar TS4500CLI.jar -ip 10.1.121.99 -u Redbooks -p RedB00ks --view
RolePermissions Redbooks > rolenew.txt
                              Action,   Access Level
                                 AIT,     No Access
                          Cartridges,        Modify
      Cartridges by Logical Library,        Modify
               Cleaning Cartridges,        Modify
                     Debug Messages,     No Access
                              Drives,        Modify
          Drives by Logical Library,        Modify
                Email Notifications,     Read Only
                   Email Recipients,     Read Only
                Encryption Internal,     Read Only
             Encryption Key Manager,        Modify
                     Ethernet Ports,        Modify
                              Events,        Modify
                Library Information,        Modify
                  Licensed Functions,        Modify
                   Logical Libraries,        Modify
            Management GUI Behavior,        Modify
                     Master Console,     No Access
                     Password Rules,        Modify
                   Position Control,     No Access
              Remote Authentication,        Modify
                               Roles,        Modify
                   SNMP Destinations,        Modify
```

```
              SNMP Notifications,         Modify
                     Scan Speed,      No Access
            Secure Socket Layer,         Modify
                   Service Port,      No Access
          Syslogs Notifications,         Modify
                         System,         Modify
           System Date and Time,         Modify
                          Tasks,         Modify
                          Users,         Modify
                  VOLSER Ranges,         Modify
VOLSER Ranges by Logical Library,         Modify
```

For more information about the TS4500, see Chapter 5, “Command-line interface” on page 325.

4.5 Cartridges menu

Use the TS4500 management GUI pages that are available under the Cartridges icon to view all cartridges, view cartridges by logical library, assign and move cartridges, and modify logical libraries, as shown in Figure 4-64.

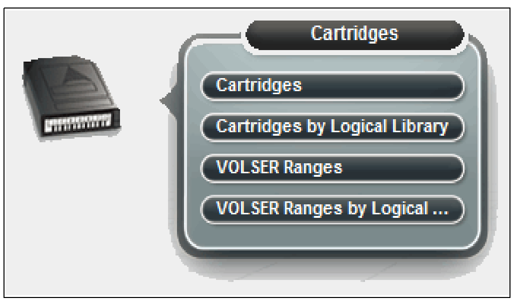

Figure 4-64 Cartridges menu

4.5.1 Cartridges

Use the Cartridges page to move cartridges to a different location, remove cartridges to an I/O station, assign cartridges to a logical library, or unassign cartridges from a logical library.

The Cartridges menu has the following functions:

- Searching for cartridges
- Viewing cleaning cartridge status
- Viewing cartridge states
- Assigning or unassigning cartridges
- Fixing a tape with an unknown volume serial number (VOLSER)
- Selecting priority for media verification
- Moving cartridges
- Exporting the mount history
- Performing other cartridge tasks from the CLI

Figure 4-65 shows the main Cartridges window.

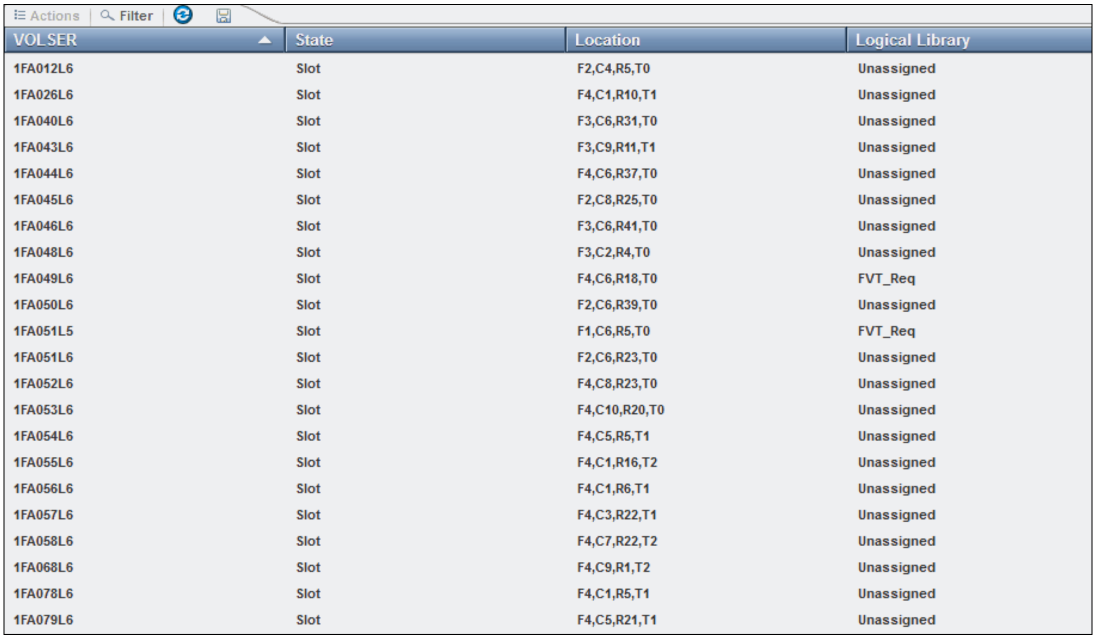

VOLSER	State	Location	Logical Library
1FA012L6	Slot	F2,C4,R5,T0	Unassigned
1FA026L6	Slot	F4,C1,R10,T1	Unassigned
1FA040L6	Slot	F3,C6,R31,T0	Unassigned
1FA043L6	Slot	F3,C9,R11,T1	Unassigned
1FA044L6	Slot	F4,C6,R37,T0	Unassigned
1FA045L6	Slot	F2,C8,R25,T0	Unassigned
1FA046L6	Slot	F3,C6,R41,T0	Unassigned
1FA048L6	Slot	F3,C2,R4,T0	Unassigned
1FA049L6	Slot	F4,C6,R18,T0	FVT_Req
1FA050L6	Slot	F2,C6,R39,T0	Unassigned
1FA051L5	Slot	F1,C6,R5,T0	FVT_Req
1FA051L6	Slot	F2,C6,R23,T0	Unassigned
1FA052L6	Slot	F4,C8,R23,T0	Unassigned
1FA053L6	Slot	F4,C10,R20,T0	Unassigned
1FA054L6	Slot	F4,C5,R5,T1	Unassigned
1FA055L6	Slot	F4,C1,R16,T2	Unassigned
1FA056L6	Slot	F4,C1,R6,T1	Unassigned
1FA057L6	Slot	F4,C3,R22,T1	Unassigned
1FA058L6	Slot	F4,C7,R22,T2	Unassigned
1FA068L6	Slot	F4,C9,R1,T2	Unassigned
1FA078L6	Slot	F4,C1,R5,T1	Unassigned
1FA079L6	Slot	F4,C5,R21,T1	Unassigned

Figure 4-65 Cartridges

Searching for cartridges

You can use the Cartridges table to find specific cartridges and view their state (slot, I/O slot, drive, or gripper). Sort the table to locate specific tape cartridges. You can sort by VOLSER, state, location, logical library, element address, or most recent use. Go to the Cartridges by Logical Library page to see the cartridges, which are displayed by their logical library.

Search for specific cartridges by using the filter. Click **Filter** and enter a value in the search field or click the **Advanced Filter** icon to the right of the search field to enter extended search criteria, as shown in Figure 4-66. Enter substrings if the entire value is unknown. Wildcard characters are not supported. Instead of wildcard characters, you can add substrings to the search by clicking the plus sign (**+**) at the end of each criterion. (Search values are not case-sensitive.)

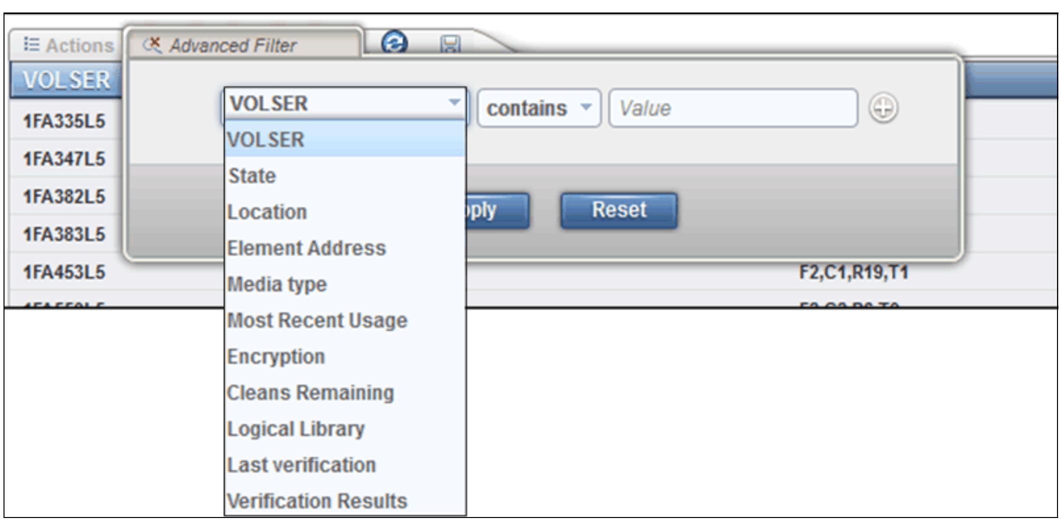

Figure 4-66 Advanced Filter option

The Cartridge database can be saved to a `.csv` file from the save option. It is also possible to save the list of cartridges that results from a search. To export the cartridge data, select the **Export Data** or **Save** icon (diskette), and then, select **Export Table Data**, as shown in Figure 4-67.

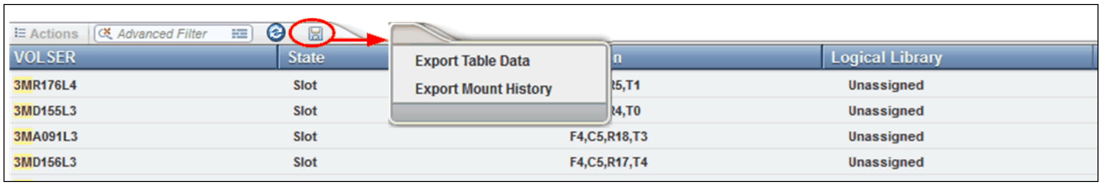

Figure 4-67 Save list

Viewing cleaning cartridge status

IBM supplies a cleaning cartridge with the first frame of each media type in a library. The library uses the cleaning cartridge to automatically clean tape drives, as needed, to maintain the efficiency of the tape drives. Each cleaning cartridge can be used up to 50 times.

This option is used to view how many cleanings remain on a cleaning cartridge. The number of remaining cleanings is displayed in the Cleans Remaining column of the Cartridges table. This column is hidden, by default.

Right-click the table header to select the Cleans Remaining column. To see all cleaning media in the TS4500, sort the column by Cleans Remaining by using the arrow, as shown in Figure 4-68.

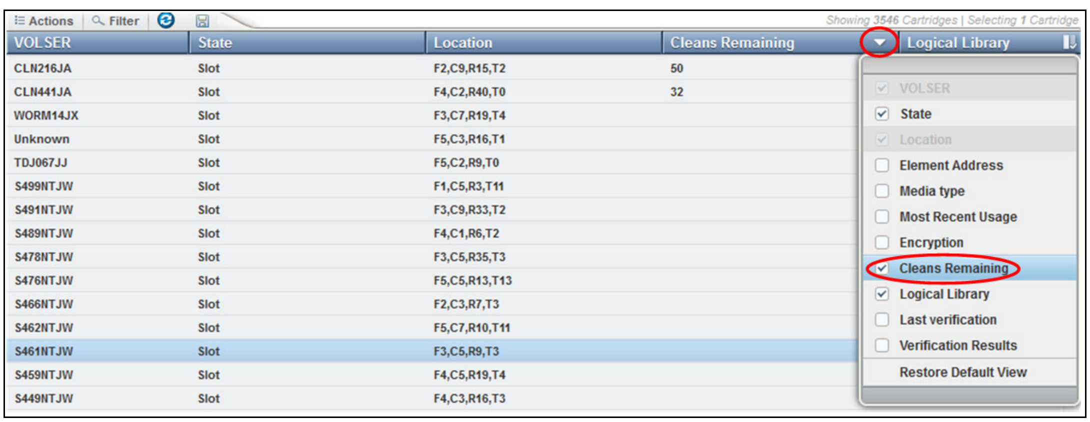

Figure 4-68 Remaining cleanings

You can also view how many cleaning cycles remain on all cleaning cartridges in a logical library on the Cartridges by Logical Library page.

You can optionally eject expired cleaning cartridges automatically. For more information about enabling this feature, see “Advanced” on page 194.

Note: Cleaning cartridges are shared among logical libraries.

Cartridge state

The State column in the Cartridges table lists the current state for each cartridge. The following states are possible:

- Drive: A tape cartridge is in a tape drive.
- Gripper: A tape cartridge is in the gripper of the cartridge accessor.
- I/O Slot (Import Queued): A tape cartridge is in an I/O slot, but it is queued to move to a storage slot.
- I/O Slot: A tape cartridge was moved to the I/O station. This state is cleared if the cartridge is moved by the operator to any other location, including a different I/O station slot.
- Slot: A tape cartridge is in a storage slot.
- Slot (Assignment Pending): A tape cartridge is in the Unassigned logical library because an empty import/export element (IEE) address is not currently available for assignment.
- Slot (Export Queued): A tape cartridge was queued to move from a slot to the I/O station by one of the following methods:
  - The host issues a `move` command from the storage element to the IEE.
  - An operator selects a Move To I/O Slot operation from the management GUI.
- Slot (Virtual I/O): A tape cartridge was imported (physically moved) into the slot, but it is still in the virtual I/O (VIO) element address space and it is not queued for export.

Assigning or unassigning cartridges

A bar code label with a volume serial (VOLSER) number is affixed to each cartridge. Cartridges are assigned to a logical library with VOLSER ranges that are defined when the logical library is created. If the VOLSER range that is assigned to a logical library matches the VOLSER of a cartridge, that cartridge is assigned to the logical library. If the VOLSER of a newly inserted cartridge does not match the VOLSER range of a logical library, the cartridge is available to import into any logical library of the same media type.

The assignment is then determined by the first application to import the cartridge. VOLSER ranges were called the *Cartridge Assignment Policy* with the TS3500 tape library.

You can modify the cartridges that are assigned to a logical library by modifying the VOLSER range of the logical library from the VOLSER Ranges by Logical Library page.

Assigning a cartridge

If an individual cartridge is outside of a VOLSER range, you can use the Assign action to assign that cartridge to a specific logical library. However, creating VOLSER ranges is the preferred method for assigning cartridges to logical libraries.

To assign or reassign cartridges to a logical library, highlight one or more cartridges from the same logical library and then, select **Actions** → **Assign**, as shown in Figure 4-69.

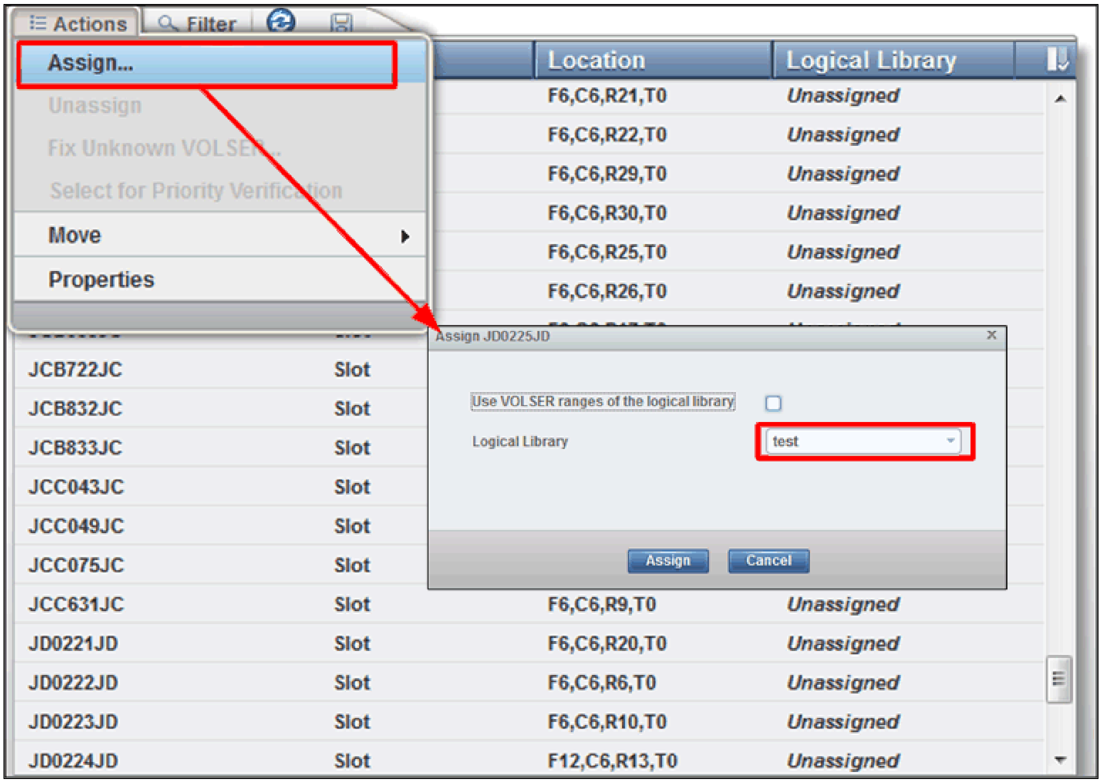

Figure 4-69 Assign selection

You can bulk assign all available or unassigned cartridges from the Cartridges by Logical Library page by highlighting the Available or Unassigned cartridge rows and selecting **Actions** → **Assign All Cartridges**.

Note: If the new cartridge assignments exceed the maximum number of cartridges that is allowed by a logical library, any other cartridges are unassigned. Go to the Logical Libraries page to modify the maximum number of cartridges that is allowed, and then attempt to assign the cartridge again.

Unassigning cartridges

To unassign one or more cartridges, select the cartridge, and click **Actions** → **Unassign**. Select **Yes** in the pop-window, as shown in Figure 4-70. However, as shown in the warning message, this action might result in synchronization problems with the application. To avoid this scenario, the preferred method is to use the application to eject cartridges.

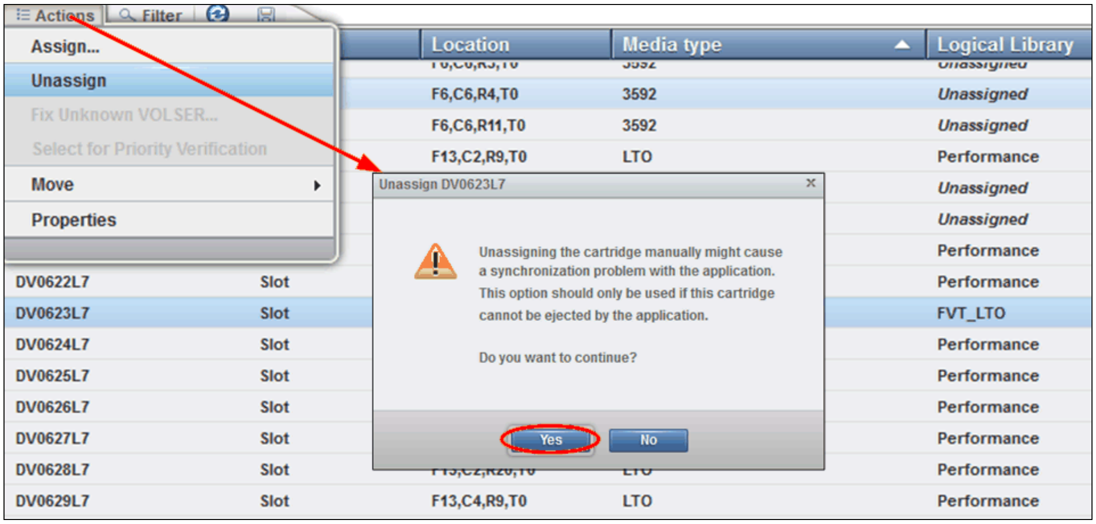

Figure 4-70 Selecting the Unassign option

Fixing a tape with an unknown VOLSER

Many tape management applications use Standard Label tape processing. The VOLSER number on the bar code label must match the VOLSER that is written to the tape to maintain compatibility with this type of processing. If a tape's VOLSER is unknown, you cannot use it until a VOLSER is established.

If a tape has an unknown VOLSER, select **Actions** → **Fix Unknown VOLSER** to apply a new VOLSER number to the tape. Enter the first six characters of the volume serial number and then, add the last two characters that specify the cartridge type, taking note of the warning message, as shown in Figure 4-71.

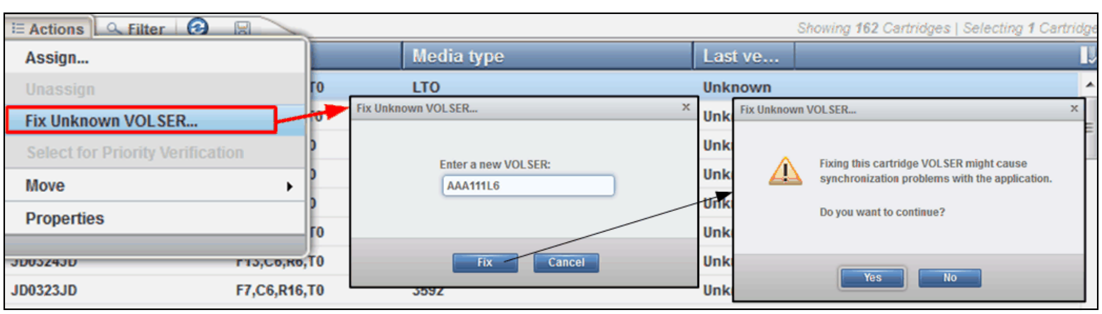

Figure 4-71 Fix Unknown VOLSER

Table 4-6 lists the possible cartridge type values.

Table 4-6 Cartridge types

Media	Cartridge type
LTO media	L8, LY, L7, LX, L6, LW, L5, LV, L4, LU, L3, and LT
3592 media	JE, JV, JM, JD, JZ, JL, JC, JY, JB, JX, and JK

VOLSER

> **Note:** This action might result in synchronization problems with the application. An audit type of action for your application might be required.

Select for Priority Verification option

After you set up the verification function, you can select one tape volume to verify. Select the media from the Cartridges menu and select the VOLSER to verify. In the Actions menu, click **Select for Priority Verification**, as shown in Figure 4-72. For information about media verification, see "Media verification" on page 258.

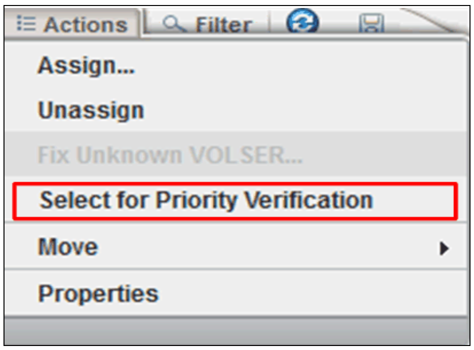

Figure 4-72 Select for Priority Verification option

Moving cartridges

At times, you might want to instruct the library to move a specific tape cartridge. For example, if a single host controls the library and the host fails during an operation, you can use the Actions menu on the Cartridges page to move one or more cartridges to continue the operation, as shown in Figure 4-73.

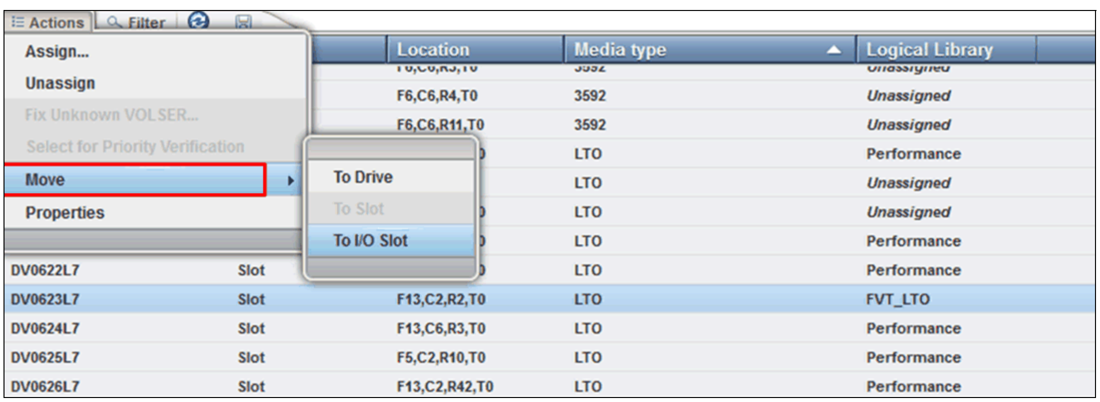

Figure 4-73 Move cartridge

You can select **Actions** → **Move** → **To Drive**, **Actions** → **Move** → **To Slot**, or **Actions** → **Move** → **To I/O Slot**, depending on where you want to move the cartridge.

> **Note:** The use of the Cartridges page to move cartridges might result in synchronization problems with the application. To avoid this scenario, the preferred method is to use the application to move cartridges.

Exporting the mount history

Click the **Save** icon (see Figure 4-67 on page 240) and click **Export Mount History** to download the Mount History report. This report is a comma-separated value (`.csv`) file that contains a history of the last 250 unmounted cartridges and information about the tape drives. An example of a portion of this file is shown in Figure 4-74.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	Mount History													
2							Mount	>>>	>>>	>>>	>>>	>>>	Life	>>>
3							Media	>>>	Host	>>>	Drive	Tape Alert	Mounts	WRetries
4							Write	Read	Write	Read	Residency	Media	Media	Media
5	Date and Time	VolSer	Frame	Drive	LogLib	EAddr	MBs	MBs	MBs	MBs	Minutes	_1-64	Count	Count
6	2014 05 15 18:56	_D040NTJR	3	11	_LL4Beta	258	0	0	0	0	0	0	0	0
7	2014 05 15 18:56	_JJS249JJ	2	14	_LL4Beta	259	0	0	0	0	0	0	0	0
8	2014 05 15 18:57	_9T0028JA	3	10	_LL4Beta	257	0	0	0	0	0	0	0	0
9	2014 05 15 18:59	_9T0028JA	2	14	_LL4Beta	259	0	0	0	0	0	0	0	0
10	2014 05 15 20:01	_D040NTJR	3	10	_LL4Beta	257	0	0	0	0	0	0	0	0
11	2014 05 15 20:01	_JJS249JJ	3	11	_LL4Beta	258	0	0	0	0	0	0	0	0
12	2014 05 15 21:06	_A00008	3	10	_LL4Beta	257	0	0	0	0	0	0	0	0
13	2014 05 15 21:07	_D040NTJR	3	10	_LL4Beta	257	0	0	0	0	0	0	0	0
14	2014 05 12 16:23	_313034JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
15	2014 05 12 16:26	_313005JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
16	2014 05 12 16:33	_313099JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
17	2014 05 12 17:14	_313034JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
18	2014 05 12 17:21	_313099JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
19	2014 05 12 17:32	_313034JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
20	2014 05 12 17:33	_313005JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
21	2014 05 12 17:36	_313034JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
22	2014 05 12 17:37	_313099JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0
23	2014 05 12 17:41	_313005JB	3	15	_FVT1_Tal	257	0	0	0	0	0	0	0	0

Figure 4-74 Mount History

The file name includes the date that the file was downloaded and the library name (`LIBNAME_YYYYMMDD_mount_history.csv`).

The following information is included in the mount history report:

- Date and Time: The date and time that the cartridge was mounted in the drive. The format is `yyyy mm dd hh:nn.ss`, where `yyyy` is the year, `mm` is the month, `dd` is the day, `hh` is the hour, `nn` is the minute, and `ss` is the second.
- VolSer: The volume serial number of the cartridge (also known as the *VOLSER*). The VOLSER is a unique identifier. Ignore the underscore that precedes the VOLSER.
- Frame: The number of the library frame from which the cartridge was unmounted. Beginning with the base frame, frames are numbered 1 - 4, from left to right.
- Drive: The number of the drive from which the cartridge was unmounted. Drives are numbered 1 - 16.
- LogLib: The name of the logical library to which the cartridge was assigned.
- EAddr: The element address from which the cartridge was unmounted.
- Mount Tape Alert Media: The number of the most recent TapeAlert flag that was received by this drive and that referred to this cartridge.
- Life Mounts Media: The number of times that the cartridge was mounted to a drive since it was manufactured.
- Life WRetries Media: During the life of the cartridge, the number of errors that occurred when drives tried write operations again.
- Life WPerms Media: During the life of the cartridge, the number of permanent, unrecoverable errors that occurred when drives performed write operations.
- Life RRetries Media: During the life of the cartridge, the number of errors that occurred when drives tried read operations again.
- Life RPerms Media: During the life of the cartridge, the number of permanent, unrecoverable errors that occurred when drives performed read operations.

- Mount Rating Drive: The overall measure of the condition of the drive. `X'00'` is unknown. The value ranges from `X'01'` (best) to `X'FF'` (worst). This rating is the efficiency of the drive.
- Mount Rating Media: The overall measure of the condition of the cartridge that is mounted. `X'00'` is unknown. The value ranges from `X'01'` (best) to `X'FF'` (worst). This rating is the efficiency of the cartridge.
- Mount Rating Ports: The overall measure of the condition of the interface to the host server. `X'00'` is unknown. The value ranges from `X'01'` (best) to `X'FF'` (worst). This rating is the efficiency of the interface.
- Mount Rating Port0: The overall measure of the condition of the Port 0 interface to the host server. `X'00'` is unknown. The value ranges from `X'01'` (best) to `X'FF'` (worst). This rating is the efficiency of the Port 0 interface.
- Mount Rating Port1: The overall measure of the condition of the Port 1 interface to the host server. `X'00'` is unknown. The value ranges from `X'01'` (best) to `X'FF'` (worst). This rating is the efficiency of the Port 1 interface.
- Mount Rating Rsvd: Reserved for the library interface.
- Mount Write Perf: The ratio of performance write commands in relation to all write-type commands. This value is a measure of the efficiency of write performance. The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Write ERPs: The measure of how the data rate performance affects the error-recovery procedures (ERPs) on write operations. The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Write Burst: For write operations, the measure of the comparison between the window tape buffer rate to the average rate. The *window rate* is the amount of data that is moved divided by the time when the data is ready in the mode (when data can be moved, but it is not). The average rate is the amount of data that is moved divided by the overall time in the mode (including setup, overhead, and so on). The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Write Buffer: The average tape-buffer efficiency on write operations. This value is the *streaming write efficiency*. The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Read Perf: The ratio of performance read commands in relation to all read-type commands. The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Read ERPs: The measure of how the data rate performance affects the error-recovery procedures (ERPs) on read operations. The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Read Burst: For read operations, the measure of the comparison between the window tape buffer rate to the average rate. The window rate is the amount of data that is moved divided by the time when it is ready in the mode (when data can be moved, but it is not). The average rate is the amount of data that is moved divided by the overall time in the mode (including setup, overhead, and so on). The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Read Buffer: The average tape-buffer efficiency on read operations. This value is the *streaming read efficiency*. The value is given as a percentage. A high percentage is better and a low percentage is worse.

- Mount Capacity Total: The measure of the efficiency of static capacity. This total can be viewed as the percentage of recorded media that fits into the currently recorded area in relation to how much data can ideally fit in that area. The lower the percentage, the less capacity is available (due to recording error recovery, media defects, and so on).
- Mount Capacity Writes: The measure of the efficiency of active capacity on write commands. This value can be viewed as the sum of efficiency for write operations on this mount. The value is given as a percentage. A high percentage is better and a low percentage is worse.
- Mount Capacity Control: The measure of the efficiency of active capacity on all other operations. The value is given as a percentage. A high percentage is best and a low percentage is worst.
- Crypto Status: Whether a cartridge is encrypted. Values are `1` (media contains encrypted data), `0` (media does not contain encrypted data), or a blank space if the drive was unable to determine whether the media contains encrypted data.
- Crypto Rekey: Whether a cartridge was rekeyed. Values are `1` (cartridge was rekeyed during the last mount) or `0` (cartridge was not rekeyed during the last mount).
- Mount Host Write: The number of megabytes (MB) that was written during a mount.
- Mount Host Read: The number of megabytes (MB) that was read during the mount.
- Mount Drive Residency: The number of minutes that the tape cartridge remained in the tape drive during the mount.

4.5.2 Cartridges by Logical Library

Use the Cartridges by Logical Library page to create and manage logical libraries, and assign, unassign, or move cartridges.

The Cartridges by Logical Library window has the following functions:

- Searching for cartridges
- Viewing cleaning cartridge status
- Assigning or unassigning cartridges
- Moving cartridges
- Creating and managing logical libraries

Figure 4-75 shows the main window for Cartridges by Logical Library and the available selections for the default columns that you want to view.

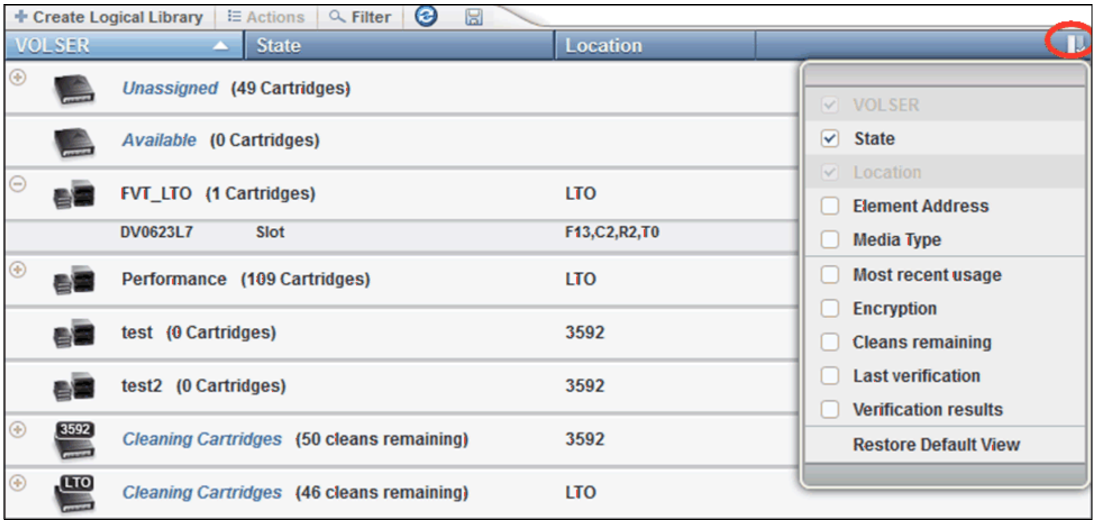

Figure 4-75 Cartridges by Logical Library

Searching for cartridges

Sort the table on the Cartridges by Logical Library page to locate specific tape cartridges within a specific logical library. You can sort by VOLSER, state, location, element address, most recent use, last verification, or verification results.

Search for specific cartridges by using the filter. Select **Filter** and enter a value in the search field. Enter substrings if the entire value is unknown. Wildcard characters are not supported. Instead of wildcard characters, you can add more substrings to the search by clicking the plus sign (+) icon at the end of each criterion. Search values are not case-sensitive.

Viewing cleaning cartridge status

IBM supplies a cleaning cartridge with the first frame of each media type in a library. The library uses the cleaning cartridge to automatically clean tape drives, as needed, to maintain the efficiency of the tape drives. Each cleaning cartridge can be used 50 times. The cleaning cartridge is shared with all logical libraries that are configured and it is not assigned to any logical library.

View how many cleanings remain on all cleaning cartridges in a logical library on the Cartridges by Logical Library table. You can also view how many cleanings remain on a specific cleaning cartridge in the Cleans Remaining column on the Cartridges page, as shown in Figure 4-76.

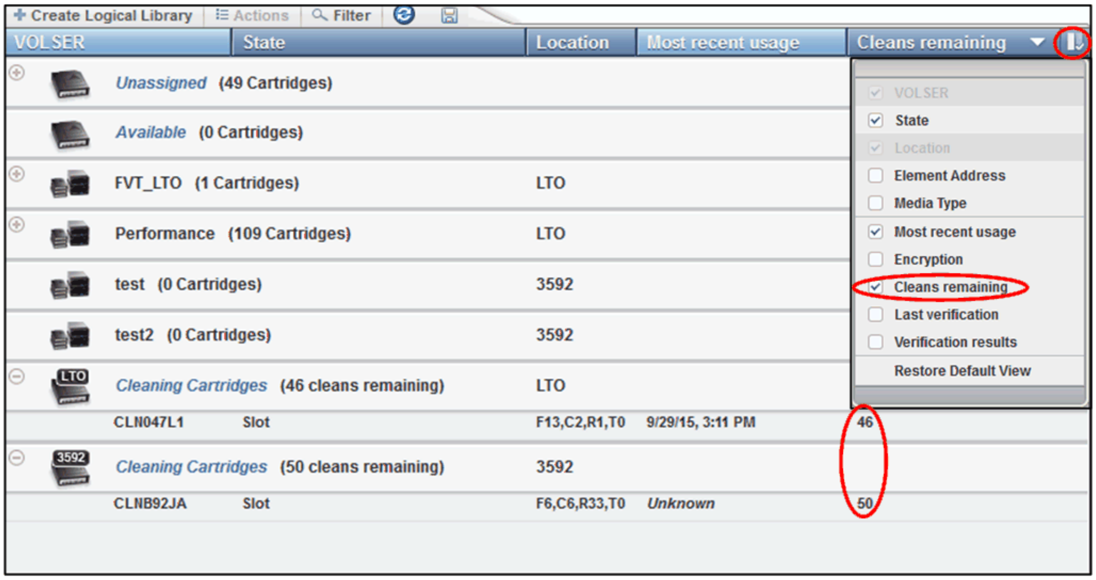

Figure 4-76 Cleaning cartridges

You can optionally automatically eject expired cleaning cartridges. To enable this feature, select **Settings** → **Library** → **Advanced**, and then, select **Automatically**, as shown in Figure 4-13 on page 194.

> **Remember:** Cleaning cartridges are shared among logical libraries.

Assigning or unassigning cartridges

A bar code label with a volume serial (VOLSER) number is affixed to each cartridge. Cartridges are assigned to a logical library with VOLSER ranges that are defined when the logical library is created. If the VOLSER range that is assigned to a logical library matches the VOLSER of a cartridge, that cartridge is assigned to the logical library.

If the VOLSER of a newly inserted cartridge does not match the VOLSER range of a logical library, the cartridge is available to import into any logical library of the same media type. The assignment is then determined by the first application to import the cartridge.

You can modify the cartridges that are assigned to a logical library by modifying the VOLSER range of the logical library from the VOLSER Ranges by Logical Library page.

Assigning

If an individual cartridge is outside of a defined VOLSER range, you can use the Assign action to assign that cartridge to a specific logical library. However, creating VOLSER ranges is the preferred method for assigning cartridges to logical libraries.

To assign or reassign cartridges to a logical library, highlight one or more cartridges from the same logical library, and then, select **Actions** → **Assign**, as shown in Figure 4-77.

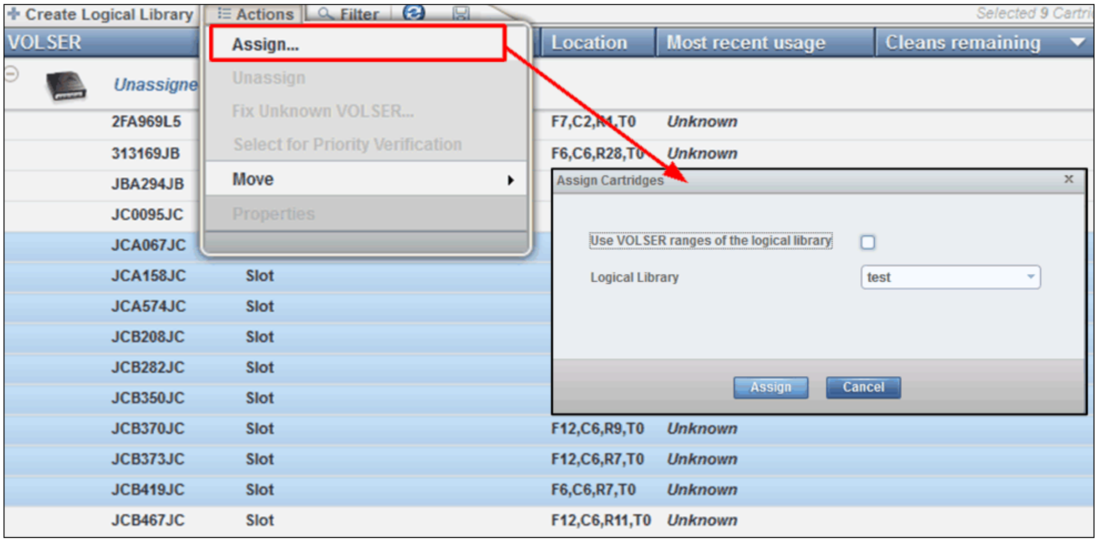

Figure 4-77 Assign Cartridges to a logical library

> **Note:** If the new cartridge assignments exceed the maximum number of cartridges that is allowed by a logical library, any other cartridges are unassigned. Go to the Logical Libraries page to modify the maximum number of cartridges that is allowed and try the cartridge assignment again.

Unassigning cartridges

Select **Actions** → **Unassign** to unassign one or more cartridges, as shown in Figure 4-78. However, this action might result in synchronization problems with the application. To avoid this scenario, the preferred method is to use the application to eject cartridges.

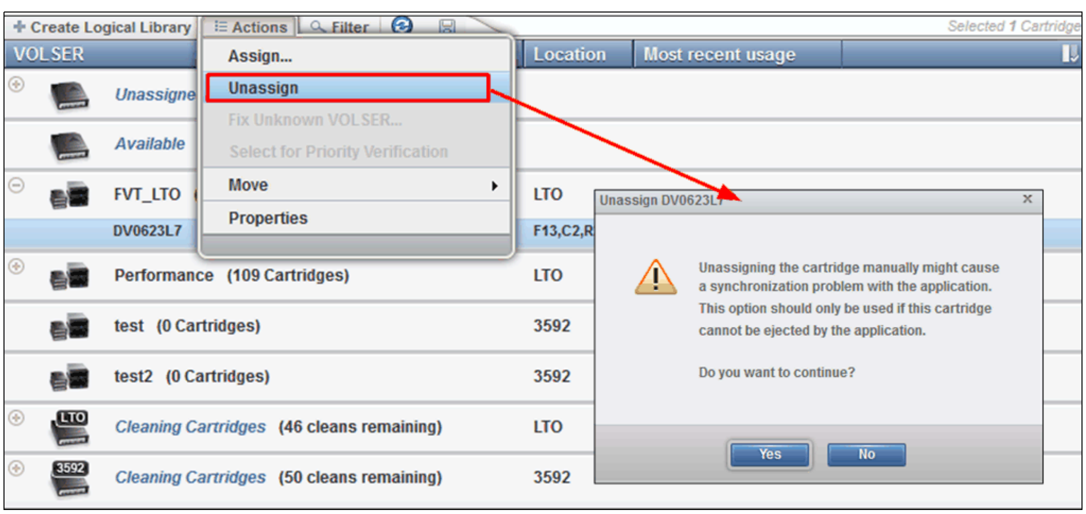

Figure 4-78 Unassigning cartridges manually

Moving cartridges

At times, you might want to instruct the library to move a specific tape cartridge. For example, if a single host controls the library and the host fails during an operation, you can use the Actions menu on the Cartridges page to move one or more cartridges to continue the operation, as shown in Figure 4-79.

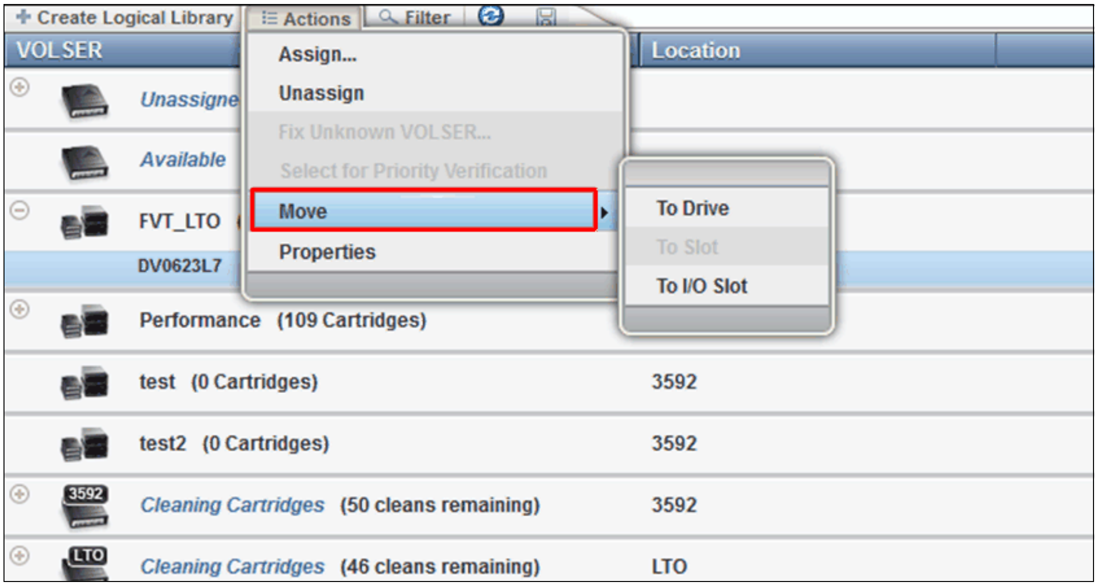

Figure 4-79 Move cartridge from the Cartridges page

Depending on where you want to move the cartridge, choose one of the following options:

- **Actions** → **Move** → **To Drive**
- **Actions** → **Move** → **To Slot**
- **Actions** → **Move** → **To I/O Slot**

Note: The use of the Cartridges page to move cartridges might result in synchronization problems with the application. To avoid this scenario, the preferred method is to use the application to move cartridges.

Creating and managing logical libraries

For more information about creating and managing logical libraries, including setting or modifying the method of encryption that is used by logical libraries, see 4.7.2, “Creating Logical Library window” on page 284.

4.5.3 VOLSER ranges

Volume serial (VOLSER) ranges are used to assign cartridges to specific logical libraries. You can view a list of all VOLSER ranges in the library, search for VOLSERs, and create, modify, and delete VOLSER ranges on the VOLSER Range page, as shown in Figure 4-80.

The following information is presented in this window:

- Using VOLSER ranges
- Creating and managing logical libraries

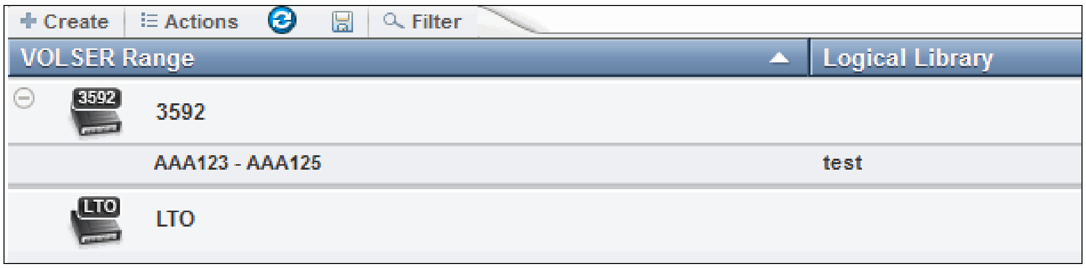

Figure 4-80 Create VOLSER Range window

When you insert a cartridge into the library and its VOLSER is within a range that is assigned to a certain logical library, the cartridge is assigned to that logical library automatically. The cartridge must be of the same media type as that logical library. For example, if you create a logical library that is called `LogicalLibrary1` for VOLSERs that range from `ABC000` to `ABC999` (a library of LTO drives) and then you insert a cartridge with VOLSER `ABC123`, the library recognizes that the VOLSER belongs in the range and assigns it to `LogicalLibrary1`.

If you insert a cartridge that is outside of any VOLSER ranges, it is available to import into any logical library of the same media type. The assignment is then determined by the first application to import the cartridge.

Within a physical TS4500 tape library, a maximum of 300 VOLSER ranges can be created among all logical libraries.

Creating or modifying a VOLSER range

To create a VOLSER range, highlight a logical library and select **Actions** → **Create VOLSER Range**. To modify a range, select **Actions** → **Modify VOLSER Range**.

Enter the start and end VOLSERs for the range and click **Create** or **Modify** (see Figure 4-81 on page 253).

Any cartridges that belong to that VOLSER range are automatically assigned, except cartridges that are being exported. If the reassignment exceeds the maximum cartridge capacity of the new or modified logical library, any excess cartridges are automatically unassigned.

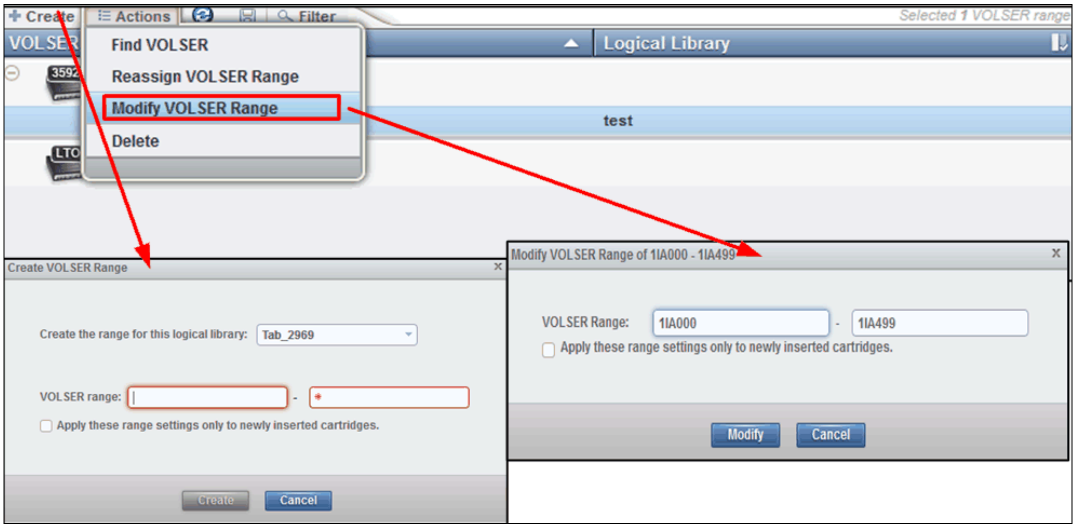

Figure 4-81 Create or modify the VOLSER range

Reassigning a VOLSER range

You can reassign a VOLSER range to a different logical library of the same media type by highlighting the range (or ranges) and by selecting **Actions** → **Reassign VOLSER Range** to display the Reassign window that is shown in Figure 4-82. Any cartridges that belong to that VOLSER range are automatically reassigned. If the reassignment exceeds the maximum cartridge capacity of the new logical library, any excess cartridges are automatically unassigned.

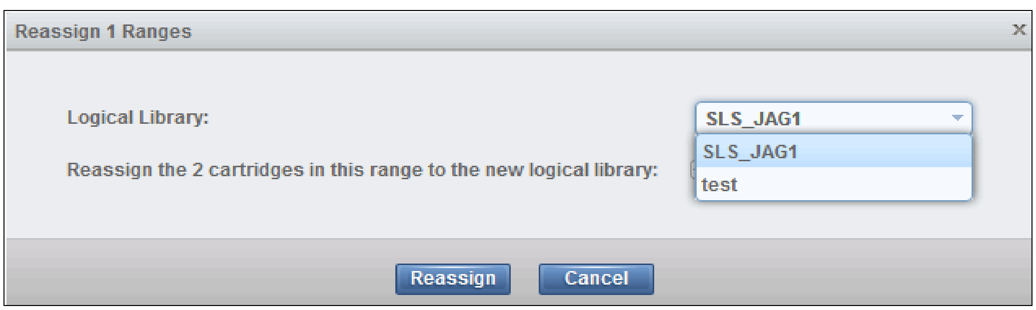

Figure 4-82 Reassign VOLSER range

Deleting a VOLSER range

To delete a VOLSER range, highlight the range and select **Actions** → **Delete VOLSER Range**. If any cartridges are assigned to the range, they are automatically unassigned.

> **Note:** Any action that causes cartridge assignments to change can result in application synchronization problems. To avoid this scenario, first use the application to eject any cartridges that are assigned to the selected VOLSER range.

4.5.4 VOLSER Range by Logical Library

Volume serial (VOLSER) ranges are used to automatically assign cartridges to specific logical libraries. You can create the ranges when you create your logical libraries or you can create, modify, and reassign the ranges on the VOLSER Range by Logical Library page. The following information is presented for using VOLSER ranges:

- Creating and managing logical libraries
- Creating, modifying, or deleting VOLSER ranges

Figure 4-83 shows the VOLSER Range by Logical Library window.

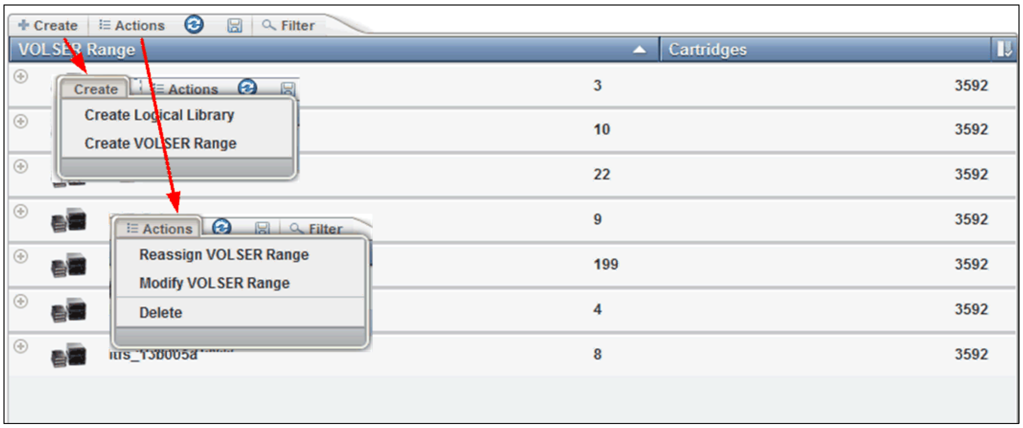

Figure 4-83 Volser Range by Logical Library window

Using VOLSER ranges

When you insert a cartridge into the library and its VOLSER is within a range that is assigned to a certain logical library, the cartridge is assigned to that logical library. The cartridge must be of the same media type as that logical library. For example, if you create a logical library that is called `LogicalLibrary1` for VOLSERs that range `ABC000` - `ABC999` (a library of LTO drives) and then you insert a cartridge with VOLSER `ABC123`, the library recognizes that VOLSER as belonging to the range and assigns it to `LogicalLibrary1`.

If you insert a cartridge that is outside of any VOLSER ranges, it is available to import into any logical library of the same media type. The assignment is then determined by the first application to import the cartridge.

Within a physical TS4500 tape library, a maximum of 300 VOLSER ranges can be created among all logical libraries.

Creating or modifying a VOLSER range

To create or modify a VOLSER range, highlight a logical library and select **Actions** → **Create VOLSER Range** or **Actions** → **Modify VOLSER Range**. Enter the start and end VOLSERs for the range and click **Create** or **Modify**. Any cartridges that belong to that VOLSER range are automatically assigned, except cartridges that are being exported. If the reassignment exceeds the maximum cartridge capacity of the new or modified logical library, any excess cartridges are automatically unassigned.

Reassign a VOLSER range

You can reassign a VOLSER range to a different logical library of the same media type by highlighting the range (or ranges) and selecting **Actions** → **Reassign VOLSER Range**. Any cartridges that belong to that VOLSER range are automatically reassigned. If the reassignment exceeds the maximum cartridge capacity of the new logical library, any excess cartridges are automatically unassigned.

Deleting a VOLSER range

To delete a VOLSER range, highlight the range and select **Actions** → **Delete VOLSER Range**. If any cartridges are assigned to the range, they are automatically unassigned.

> **Note:** Any action that causes cartridge assignments to change can result in application synchronization problems. To avoid this scenario, first use the application to eject any cartridges that are assigned to the selected VOLSER range.

4.6 Drives menu

Use the TS4500 management GUI pages under the Drives icon (see Figure 4-84) to view all drives in the library, view drives by logical library, view drive port status. The pages also can be used to perform tasks, such as assigning drives, enabling control paths, configuring drives for verification, and setting drive port configuration.

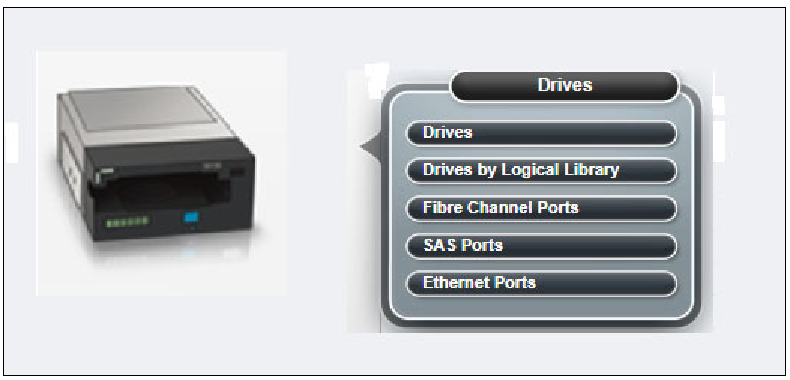

Figure 4-84 Drives menu

4.6.1 Drives

Use the Drives page to view the location and state of all drives in the library, and to enable or disable control paths, eject cartridges, reset or clean drives, or update drive firmware.

You can use the view options and display additional drive status options on the Drives page, such as port status, drive serial numbers (S/Ns), and worldwide node name (WWNN), as shown in Figure 4-85 on page 256.

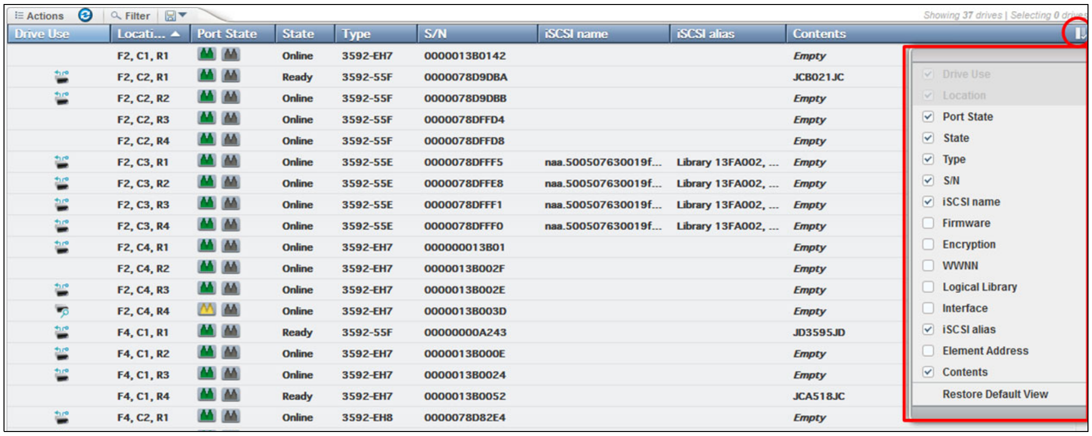

Figure 4-85 Drives page

The Drives page and the Actions menu options are shown on Figure 4-86. Depending on whether the selected drive is assigned or unassigned, you optionally can reassign or assign the drive.

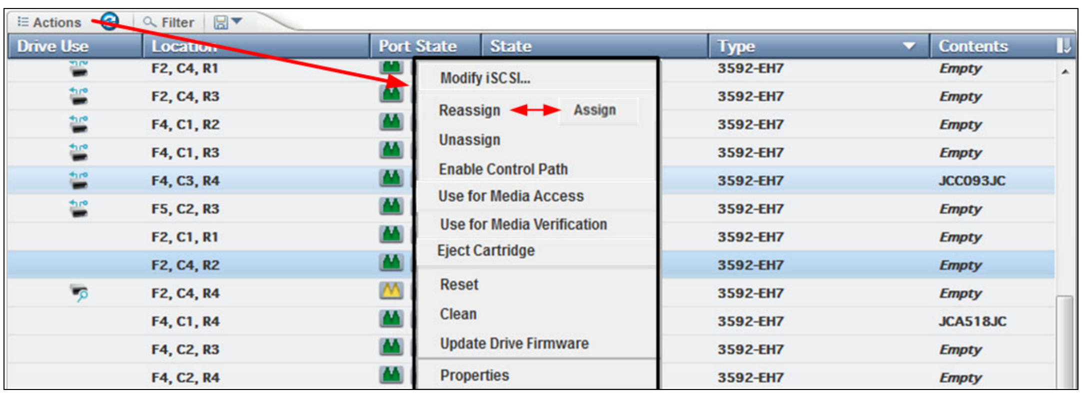

Figure 4-86 Drives page Actions menu options

Adding or removing tape drives

A tape drive is unassigned after it is installed in the tape library. All unassigned drives are assigned when a logical library is created by using the logical library presets. You also can use the Drives by Logical Library page to assign a tape drive to a logical library. For more information about the logical library configuration, see 4.7.2, “Creating Logical Library window” on page 284.

Reassigning a drive

You can reassign a drive from one logical library to another logical library by using the Reassign drive option. First, ensure that the drive is not a control path drive; then, use the Drives by Logical Library page to reassign the drive, as shown in Figure 4-87.

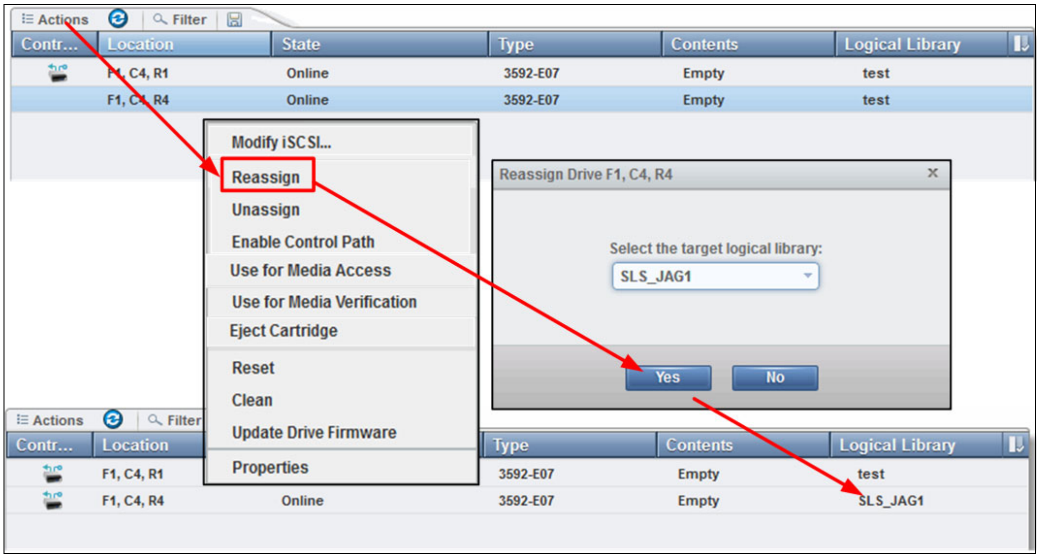

Figure 4-87 Reassign a drive

You assign a drive by using the Assign option, which is the same as reassigning the drive.

Unassigning a drive

Before a drive is physically removed from the tape library, it must be empty and unassigned. Use the Drives page to eject the cartridge and use the Drives by Logical Library page to unassign the drive, as shown in Figure 4-88.

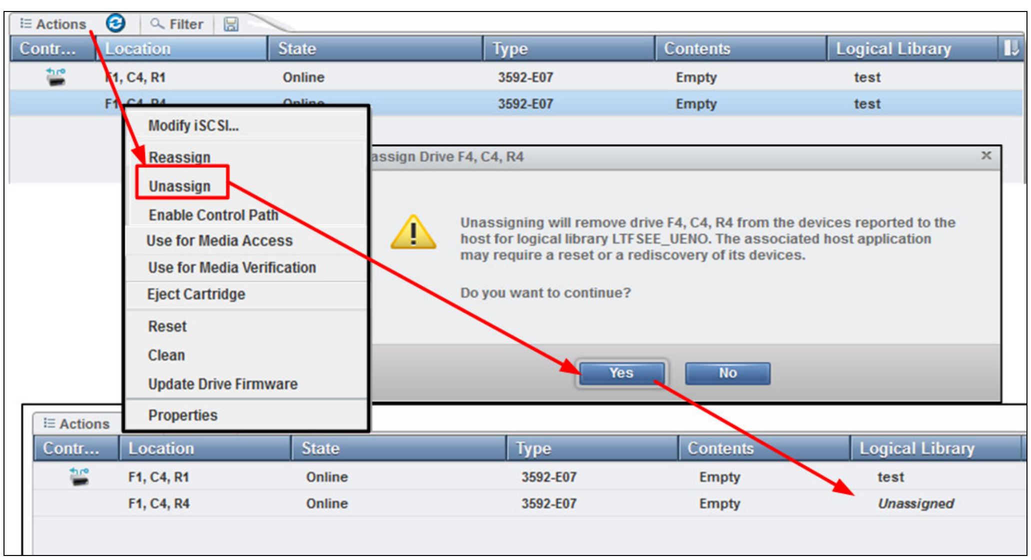

Figure 4-88 Unassign a drive

Media verification

The TS4500 supports media verification, which is fully controlled by the TS4500 at a logical library level. It is transparent to any software application that controls the library.

Use this feature to verify any media for the ability to read all of the data on the selected media within a logical library. Media verification is set up to run automatically on a set schedule.

You can set a recurring schedule for verification by days, months, or years. After you set up the media verification, you can verify a single medium before the set schedule, if required.

Cartridges that are being verified are still available to use. If the library receives a request for information about a cartridge that is being verified, the verification is suspended and the cartridge is released to fill the request. When the cartridge is no longer needed, verification is resumed.

Media verification requires that you configure dedicated drives in the logical library for this function. We recommend that you dedicate at least two drives for this task. These drives are not available to use for media access, and they are not visible to hosts.

> **Important:** Assigning drives within a logical library as media verification drives is disruptive because these drives are taken offline to any application that configured them.

For a library with multiple logical libraries, these drives can be reassigned to a different logical library to use for media verification, but this reassignment can remove the logical library verification setup.

After media verification is enabled on a logical library for the first time, all cartridges are due for verification on the same date, but because the policy period is not a deadline, it merely establishes a queue of media verification activity. The initial queue of media verification can take time, even months if a library is large.

This time depends on the number of cartridges, how full they are, and the number of media verification drives. Each cartridge has a unique last verification date, which means that all cartridges are not due for media verification at the same time.

> **Note:** We recommend that at least two drives are assigned for verification. The number of drives needs to be sized based on the number of cartridges in the logical library and how full they are.

After you run media verification on a cartridge, the next verification is based on the last verification date, plus the configured policy period of the logical library.

> **Note:** The last verification date is stored in the cartridge memory (CM) so that the policy is still applied even if a cartridge is exported and imported to the same or a different logical library (LL). The policy period is based on the logical library settings of the logical library to which the cartridge is imported.

Media verification setup

The first step is to configure a drive for media verification. The following methods are available:

- Configure an unassigned drive as the verification drive
- Configure an assigned drive as the verification drive

Configuring an unassigned drive as the verification drive

To configure an unassigned drive for media verification, from the main menu, select **Drives** → **Drives by Logical Library**, select **Unassigned drives**, and select the drive from the list to use as a verification drive.

A warning message appears to inform you that this drive will not be available for media access and it will not be visible to hosts. This drive can be used for media verification only, as shown in Figure 4-89. Repeat this task for all drives that you configure for media verification.

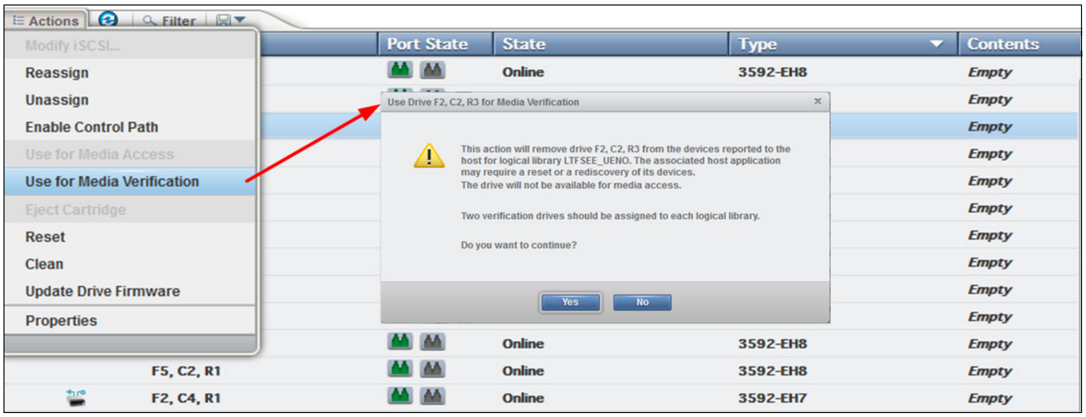

Figure 4-89 Drive configuration for media verification

After all of the required drives are set, the drives display as verification drives, as shown in Figure 4-90.

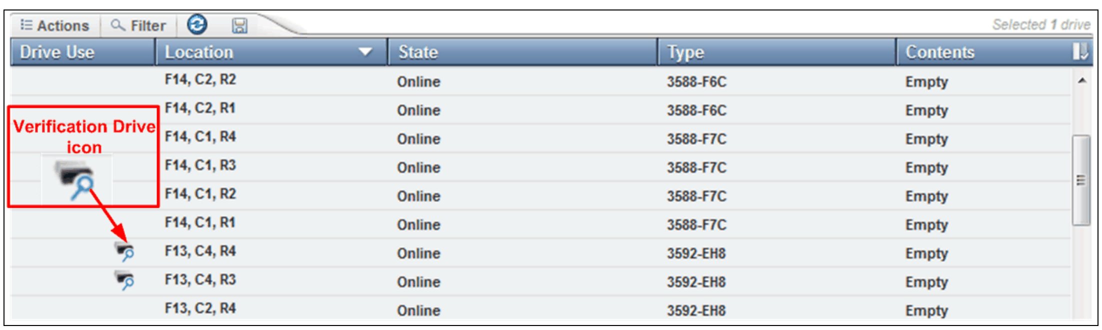

Figure 4-90 Verification drives

You can now assign these drives to a logical library that was set up for media verification. Select the drive, and then, click **Actions** → **Assign** to assign them to the logical library that is required, as shown in Figure 4-91. Repeat this task for all drives that are required to verify media.

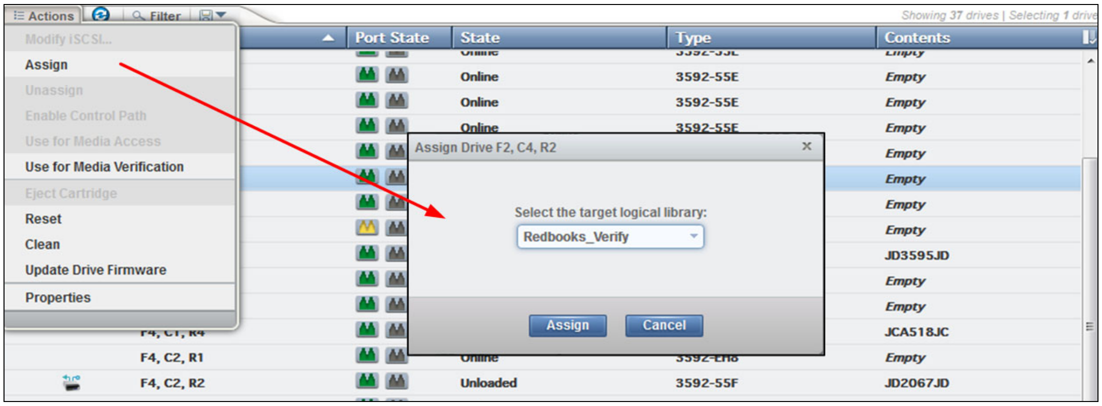

Figure 4-91 Assign verify drive

These drives are now assigned to the required logical library. They can be seen from the Drives by Logical Library menu, as shown in Figure 4-92.

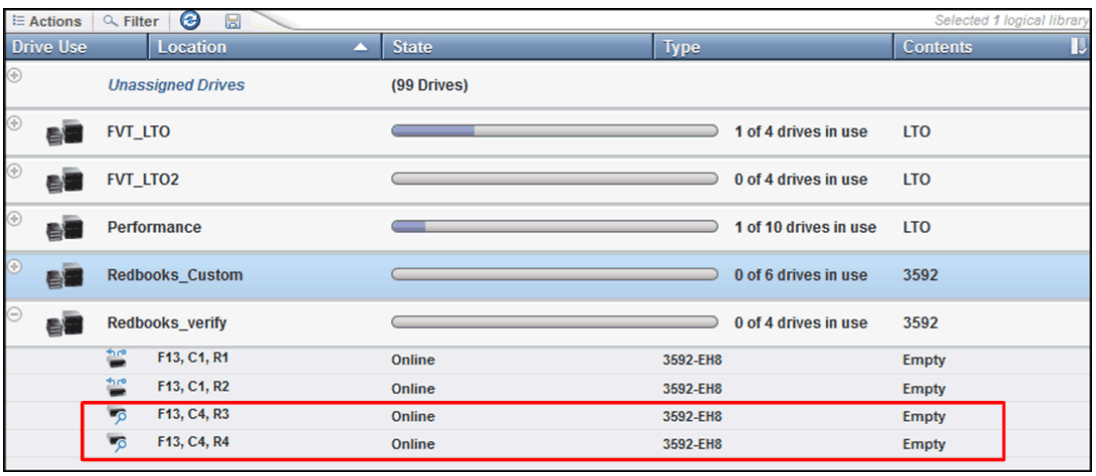

Figure 4-92 Verify media drives that are assigned to a logical library

Configuring an assigned drive as a media verification drive

From the Drives by Logical Library menu, select the logical library. Then, select the drive that you want to use for media verification. Then, click **Actions** → **Use for Media Verification**. Repeat these steps for all drives that are required for media verification. Display the drives in the logical library, and you can see that these drives are now configured for verification, as shown in Figure 4-93.

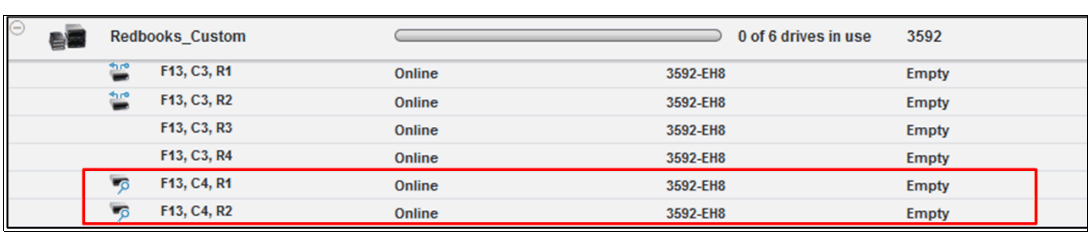

Figure 4-93 Media verification drives by logical library

Configuring media verification on a logical library

After the drives are assigned to the logical library for media verification, you can configure the logical library for media verification. Select the logical library. Then, click **Actions** → **Modify Media Verification**.

From the pop-up menu, enable **Automatic media verification** and select the frequency, first verification date, and media verification drive options to suit your requirements. You can set a recurring schedule for verification by days, months, or years.

You can select the date to start the first verification, and you can set the verification drives, as shown in Figure 4-94.

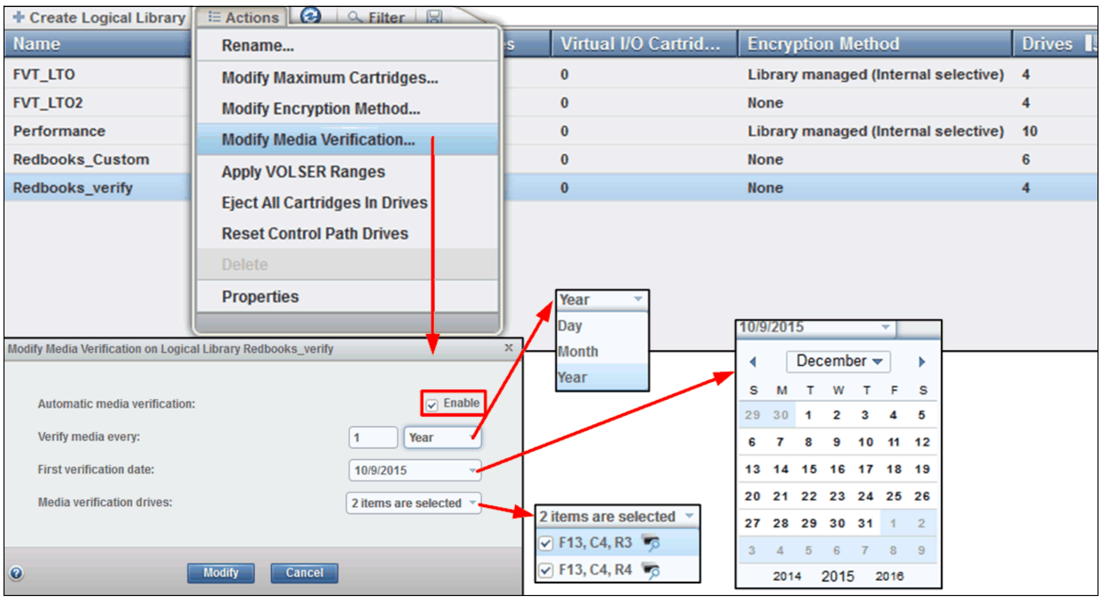

Figure 4-94 Configure media verification

This action enables media verification for the selected logical library. The verification occurs based on the policy that you set.

Verifying a single tape volume

After the verification policy is set, you can select one tape volume to verify. Select the media from the Cartridges menu, and select the VOLSER to verify. Then, click **Actions** → **Select for Priority Verification**, as shown in Figure 4-95.

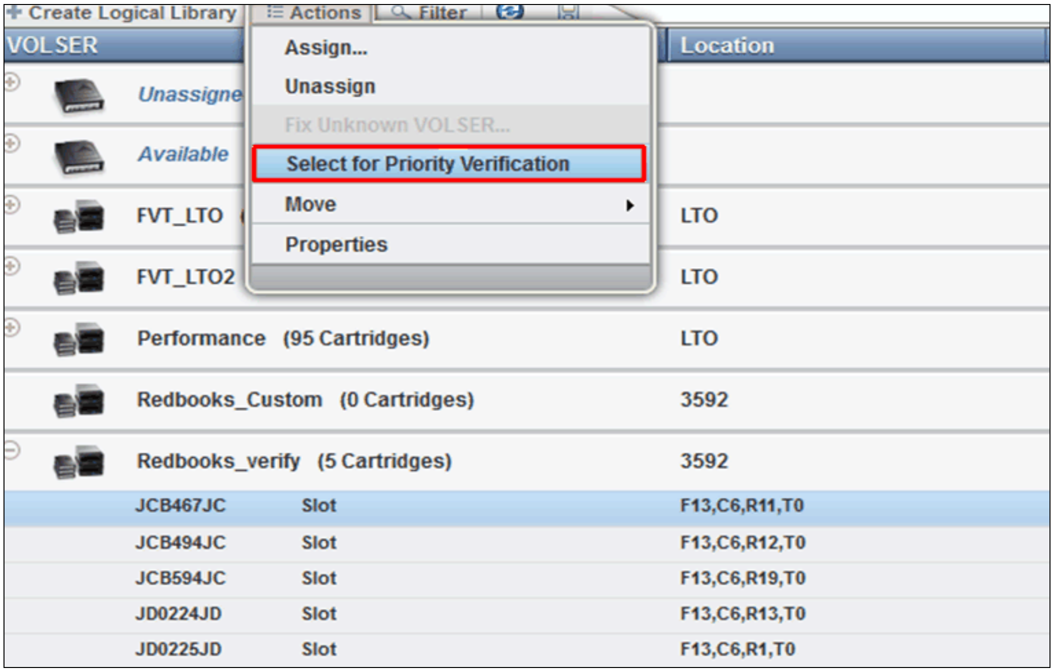

Figure 4-95 Select for Priority Verification

This selection places this tape volume next in the queue for verification. By using the Tasks menu, you can monitor and display the status of the media verification, as shown in Figure 4-96.

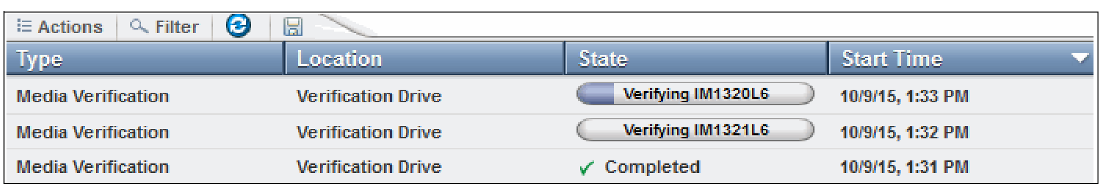

Figure 4-96 Tasks menu

You can display the status of each tape volume by selecting the tape volume and displaying its properties, as shown in Figure 4-97. The properties show the last verification date, verification result (for example, passed), and the next verification date.

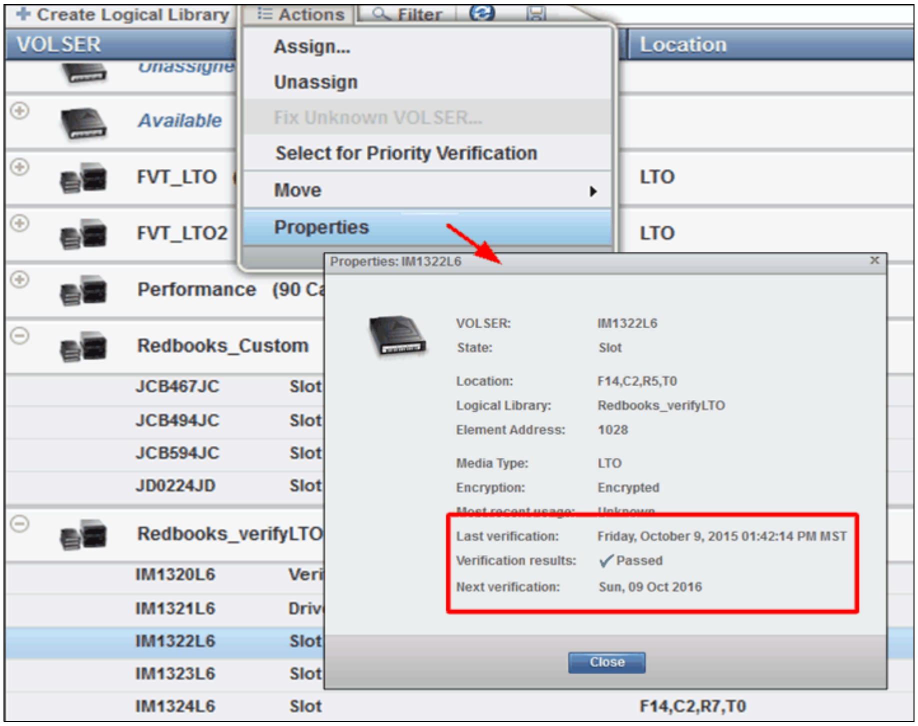

Figure 4-97 Media properties

Also, you can use the Cartridges menu to obtain the last verification date by clicking **Export Table Data** to export the table data to a spreadsheet, as shown in Figure 4-98.

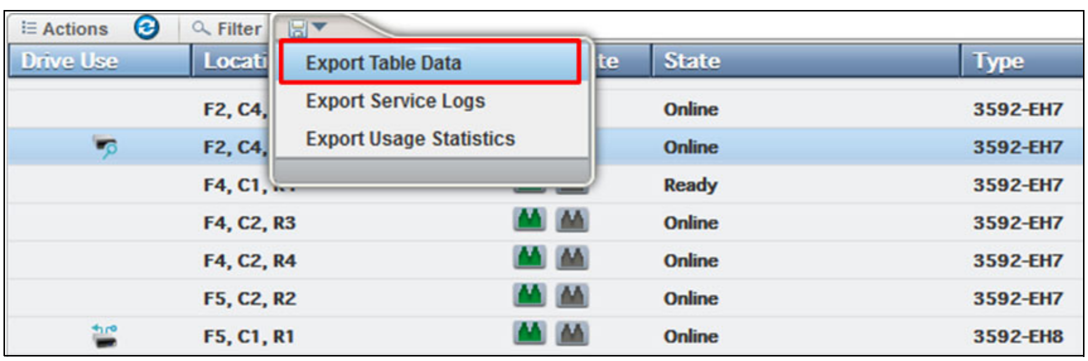

Figure 4-98 Export Table Data option

To change a drive that is assigned for media verification, select the drive in the logical library. Click **Actions** → **Use for Media Access** to set the drive for normal media access. Now, the drive is available for normal data use, as shown in Figure 4-99.

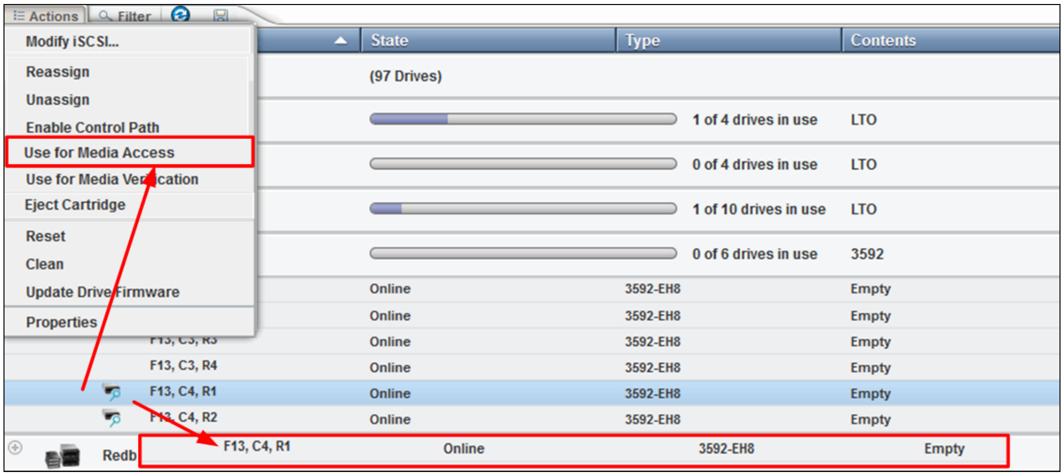

Figure 4-99 Change a verification drive to a media access drive

Control path drives

A *control path* is a logical path to the TS4500 tape library. A control path is the path for Small Computer System Interface (SCSI) Medium Changer commands that are sent by a server to control a specific logical library. The TS4500 tape library has no direct SCSI connection to a server. When a server communicates with the library, it sends the communication by way of an LTO or 3592 tape drive. The LTO or 3592 tape drive is designated as a control path.

Any drives that are enabled as control paths are identified with the icon that is shown in Figure 4-100.

Figure 4-100 Control path drive icon

To enable a control path, select the drive or drives that are to be configured, then select **Actions** → **Enable Control Path** to designate one or more drives as control path drives, as shown in Figure 4-101 on page 265.

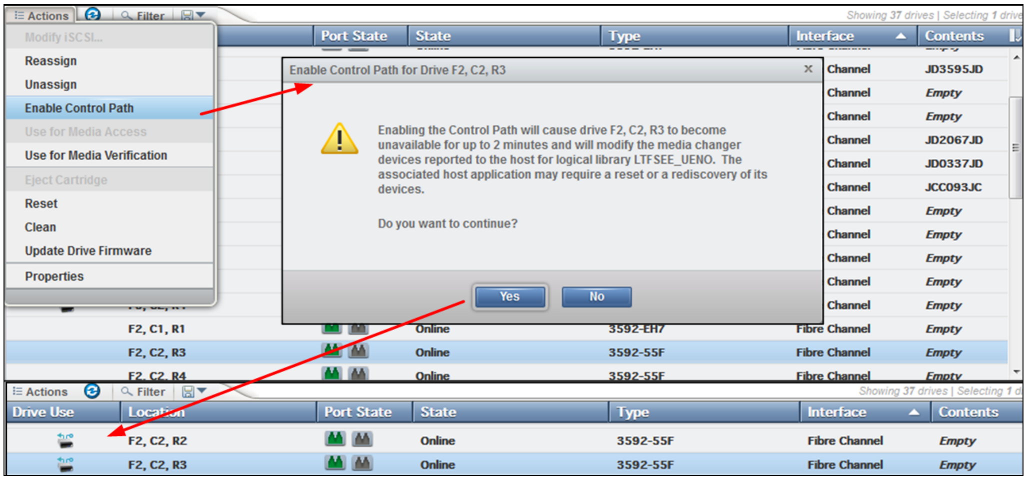

Figure 4-101 Enable Control Path

It is possible for a logical library to have no control path drives. To disable a control path, use the same method that is shown in Figure 4-101, but select **Actions** → **Disable Control Path**, as shown in Figure 4-102.

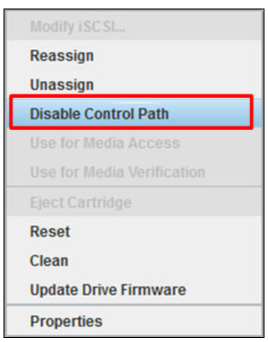

Figure 4-102 Disable Control Path

The use of multiple control paths for any single configured logical library provides the following benefits:

- Allows access to the logical library by multiple servers
- Provides communication failover so that if one control path fails, host communication for the entire library is not lost

> **Note:** Use the Path Failover feature (Feature Code 1682) to enable the host device driver to send a command again to an alternate control path in the same logical library. Configuration of Path Failover is performed on the host device driver.

Access to the logical library is on a first-come, first-served basis. Each control path for a logical library can accept commands while the library is in use by another control path.

Control path drives cannot be removed from a logical library that is performing jobs. If the logical library is not actively in use, control path drives can be removed from a logical library only after all non-control path drives are removed or a different control path drive is assigned.

iSCSI drive configuration

To configure a TS1160 model 60E or a TS1155 model 55E drive, select the drive or drives that are to be configured. Then, select **Actions** → **Modify iSCSI**. The configuration entry window opens, as shown in Figure 4-103.

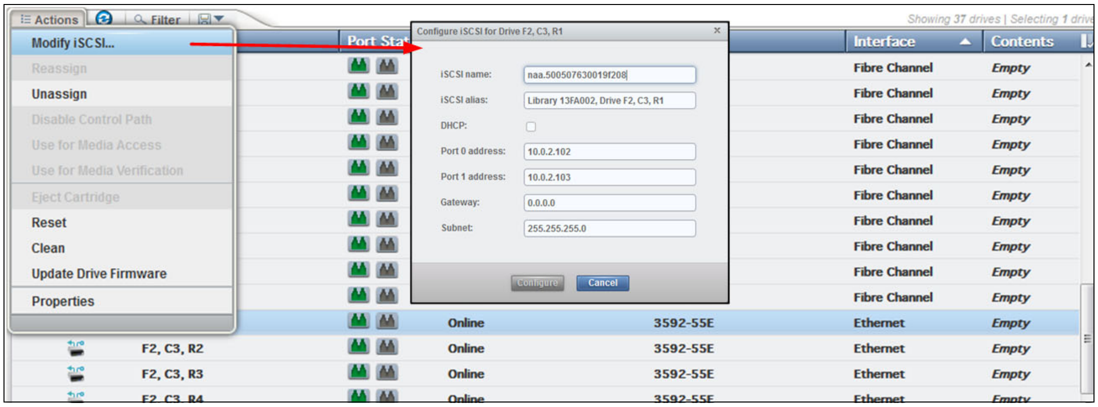

Figure 4-103 Ethernet drive configuration

In this window, you can modify or add fields, including the following examples:

- Modify the iSCSI name for this drive. This name is unique, such as an iSCSI qualified name (IQN) for each drive. The target iSCSI name has a maximum of 223 characters and has a default of `naa,driveWWNN`.
- Modify the iSCSI alias, which is a name that is given to the drive and associated to the iSCSI name. It is used to give a user-defined name to the drive, such as library serial number and drive position. This field makes managing the drive easier.
- DHCP or static IP address can be used. If DHCP is used, the IP address does not need to be added.
- Static IP address, gateway, and network mask can be added for each drive, or for a range of drives.

To configure all drives IP address information at the same time, select all of the drives that are to be configured; then, select **Actions** → **Modify iSCSI.** The window that is shown in Figure 4-104 opens. Here, you can enter the starting IP address for the drive port and all drives are configured starting at this address.

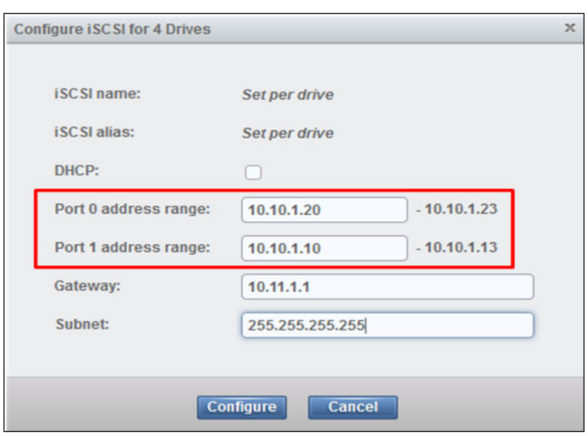

Figure 4-104 Setting range of SCSI address on TS1155 55E

CHAP can be configured for secure communications on these ports. For more information, see “iSCSI” on page 200.

To display the current drive settings, select the drive then, **Actions** → **Properties**. The current drive settings and properties are displayed, as shown in Figure 4-105.

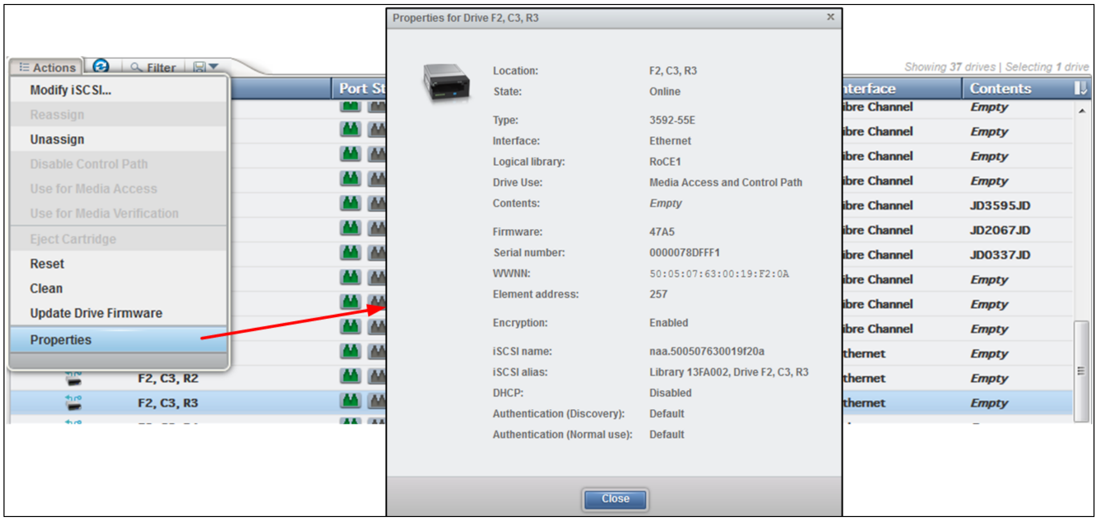

Figure 4-105 Ethernet drive property display

Resetting a tape drive

You can reset a drive by using the Drives by Logical Library page. Select the drive. Then, select **Actions** → **Reset**, as shown in Figure 4-106. The drive is unavailable for up to 4 minutes.

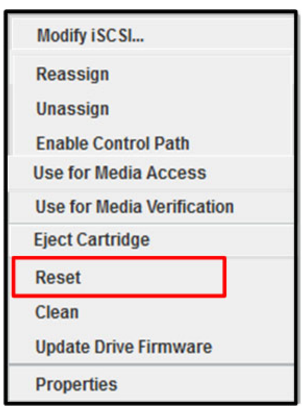

Figure 4-106 Reset the drive

Cleaning tape drives

The TS4500 tape library automatically cleans tape drives, as needed. However, you also can start a manual cleaning from the Drives page (see Figure 4-107). Select **Actions** → **Clean**.

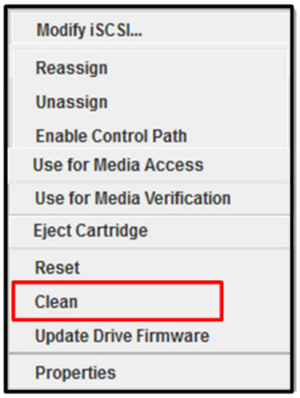

Figure 4-107 Clean a drive

A valid cleaning cartridge is required for cleaning drives. Monitor cleaning cartridge usage from the Cartridges page. For more information, see “Viewing cleaning cartridge status” on page 248.

Updating drive firmware

When you use this action (as shown in Figure 4-108), you can select one drive to update or multiple drives of the same media type simultaneously with the firmware level that was uploaded. A drive is updated even if it is running a firmware level that is later than the version of drive code that was uploaded.

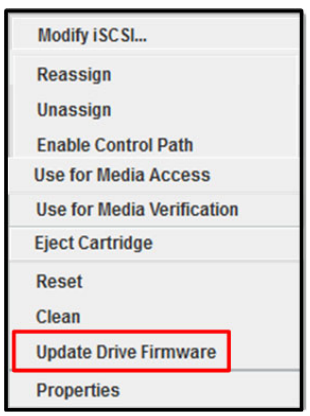

Figure 4-108 Update Drive Firmware

You can update drive firmware on one drive or multiple drives at one time. Any drives that are selected on the Drives page are updated with the firmware level that was uploaded. It is possible to install firmware levels that are earlier than the version that the drive is running.

To avoid disrupting the host, control path drives are updated when they are reset. All other drives are updated when the drive is unloaded. Control path drives must be manually reset to use the uploaded firmware level.

For more information about the drive firmware package for your workstation, see IBM Fix Central.

From the Drives page of the TS4500 management GUI, select one or more drives, and then, click **Actions** → **Update Drive Firmware**, as shown in Figure 4-108 on page 268. The File Upload window opens. Browse to the drive firmware image and click **Open**.

> **Note:** The LTO 8, 7, 6, and 5 drive firmware image file is named `LTOZ_xxxx.fcp_fh.ro` where Z is the drive type and `xxxx` is the firmware level. The TS1100 drive firmware image file is named `xxxx_xxx.fcp_fj_D.fmrz` where `xxxx_xxx` is the drive firmware level.

Monitor the status of the update in the State column and the Firmware column of the Drives table.

Select the control path drives with a state of “Reset Required” and select **Actions** → **Reset** to reset the drives.

Exporting data, logs, and statistics

Click the **Save** icon (diskette) to see the Export Data menu to export data from the Drives table, export service logs, or export drive usage statistics. This window is shown in Figure 4-109.

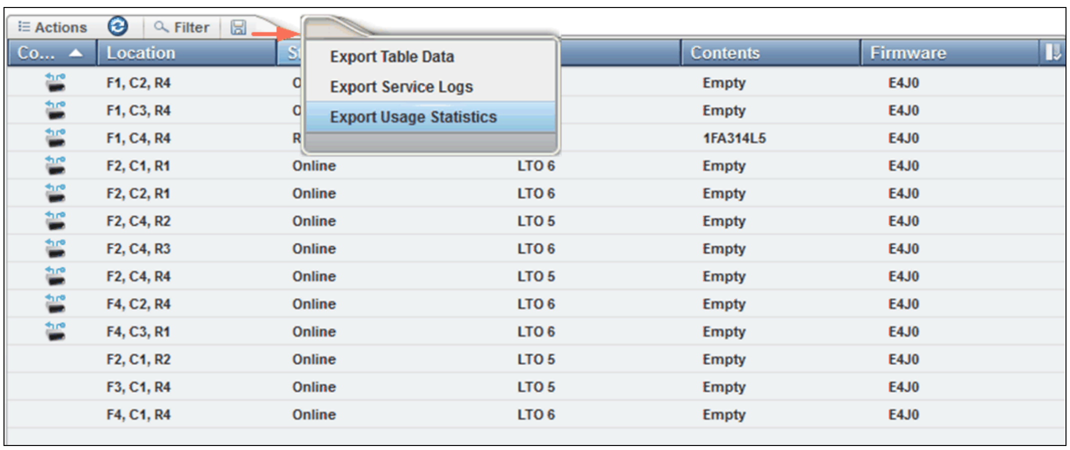

Figure 4-109 Exporting data, logs, or statistics

When you export drive usage statistics, a `LIBNAME_YYYYMMDD_drive_statistics.csv` file for all of the drives is exported. This file presents the number of loads, number of megabytes read and written, and number of cleanings for each drive location.

You can use this window to export service logs from a drive to send to IBM Support, if needed. Select the drive for which you want to export the logs, and then, select **Export Service Logs** from the Export Data menu.

4.6.2 Drives by Logical Library

Use the Drives by Logical Library page to view drive assignment and usage among all logical libraries, view cleaning cartridge status, assign drives, and manage the logical libraries and their settings. This window provides a view of the drives from a Logical Library view, as shown in Figure 4-110.

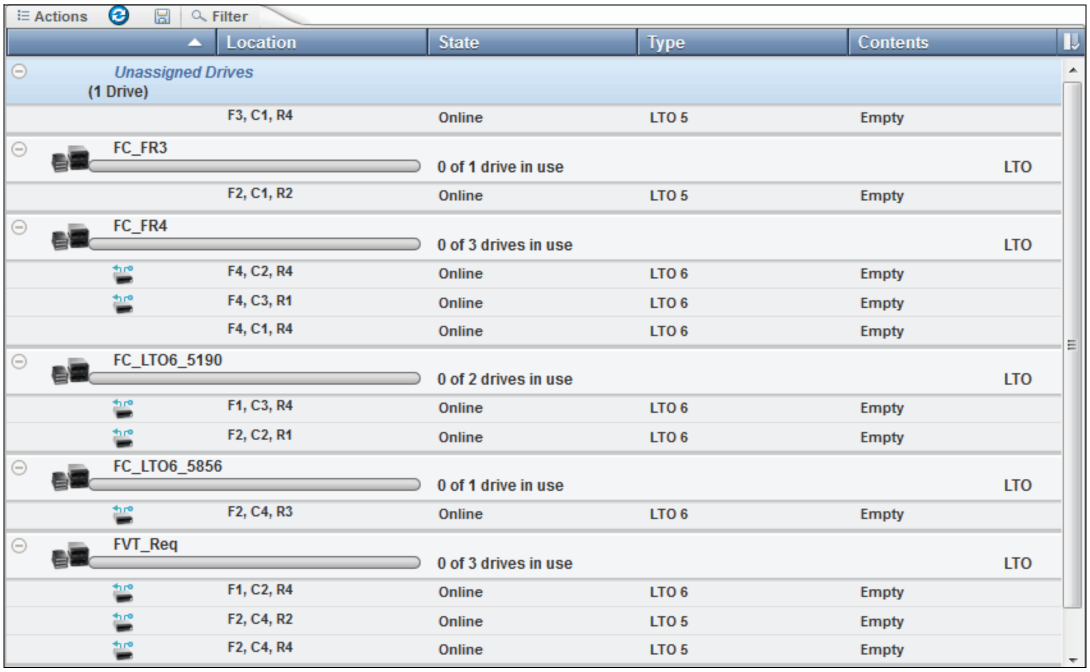

Figure 4-110 Drives by Logical Library

Monitoring drive usage

The drive usage bar indicates the current drive usage for the logical library. The number on the right side of the bar is the total number of drives that is assigned to the logical library. The bar fills as a percentage of the number of drives that is in use, as shown in Figure 4-111.

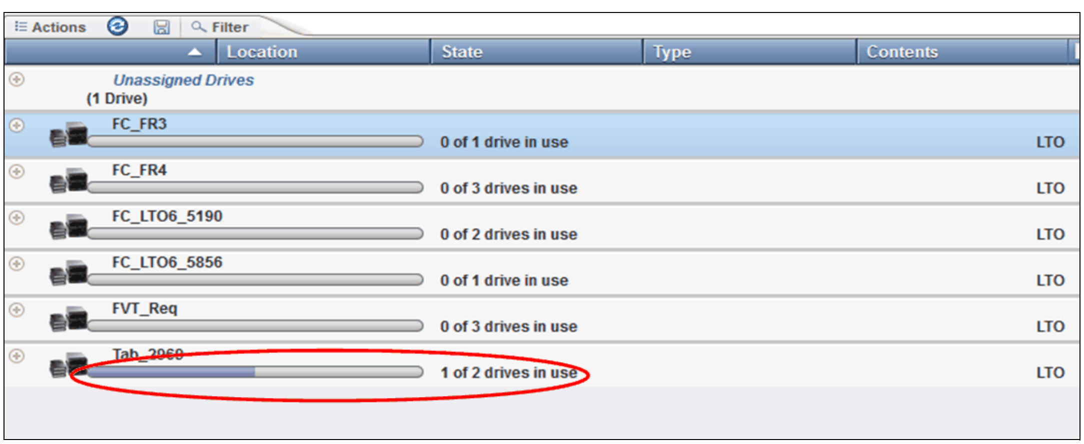

Figure 4-111 Monitor drive usage

The bar turns red if all of the drives in the logical library are being used. If all of the drives are being used frequently, the logical library might be over-allocated and you need to consider adding drives to the logical library.

The same actions can be performed on this panel as on the Drives menu. For more information about these functions, see 4.6.1, “Drives” on page 255.

4.6.3 Fibre Channel Ports

Use the Fibre Channel Ports option in the Drives menu (see Figure 4-84 on page 255) to view all of the Fibre Channel port information about drives in the library. You can view drive port status, and perform tasks, such as setting the port link speed and topology.

The selection of the Fibre Channel Ports option on the Drives menu opens a summary page of all drives and their port status, as shown in Figure 4-112.

Drive	Port	State	WWPN	Speed	Topology
F7, C3, R3	0		44:41:04:FA	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C3, R3	1		44:81:04:FA	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C3, R4	0		44:41:04:FB	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C3, R4	1		44:81:04:FB	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C4, R1	0		44:41:04:FC	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C4, R1	1		44:81:04:FC	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C4, R2	0		44:51:04:FD	Auto (8Gb/s)	Auto-L (L Port)
F7, C4, R2	1		44:91:04:FD	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C4, R3	0		44:51:04:FE	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C4, R3	1		44:91:04:FE	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C4, R4	0		44:51:04:FF	Auto (1Gb/s)	Auto-L (Auto-L)
F7, C4, R4	1		44:91:04:FF	Auto (1Gb/s)	Auto-L (Auto-L)
F8, C1, R1	0		44:51:04:00	Auto (8Gb/s)	Auto-L (L Port)
F8, C1, R1	1		44:91:04:00	Auto (1Gb/s)	Auto-L (Auto-L)
F8, C1, R2	0		44:51:04:01	Auto (8Gb/s)	Auto-L (L Port)
F8, C1, R2	1		44:91:04:01	Auto (1Gb/s)	Auto-L (Auto-L)
F8, C1, R3	0		44:51:04:02	Auto (8Gb/s)	Auto-L (L Port)

Figure 4-112 Ports summary

The port State column is sorted by state color (green, yellow, or gray).

The gray port status icon is shown in Figure 4-113.

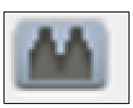

Figure 4-113 Gray port status icon

The gray port status icon indicates that no light is detected at the drive’s Fibre Channel port. One of the following conditions is true:

- A Fibre Channel cable is unplugged at the drive or at the nearest Fibre Channel device (host bus adapter, switch, or hub).
- A Fibre Channel cable is broken between the drive and the nearest Fibre Channel device (host bus adapter, switch, or hub).
- The drive or the nearest Fibre Channel device (host bus adapter, switch, or hub) is powered off or has a hardware problem.
- The drive is configured for as a media verification drive.

The yellow port status icon is shown in Figure 4-114.

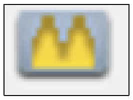

Figure 4-114 Yellow port status icon

The yellow port status icon indicates that the Fibre Channel connections between the drive and the Fibre Channel device (host bus adapter, switch, or hub) to which it is connected are good and the port is logged into the fabric. However, communication is not established to the host application (the drive might be offline to the host application).

The green port status icon is shown in Figure 4-115.

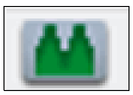

Figure 4-115 Green port status icon

The green port status icon indicates that the Fibre Channel connections are good and communication is established.

The summary display shows the Speed and Topology columns, which include both the setting and the actual value in the following format: `setting (actual)`.

To modify a port setting, select the port or ports that you want to modify and select **Actions** → **Modify Speed** or **Modify Topology** from the Ports view, as shown in Figure 4-116.

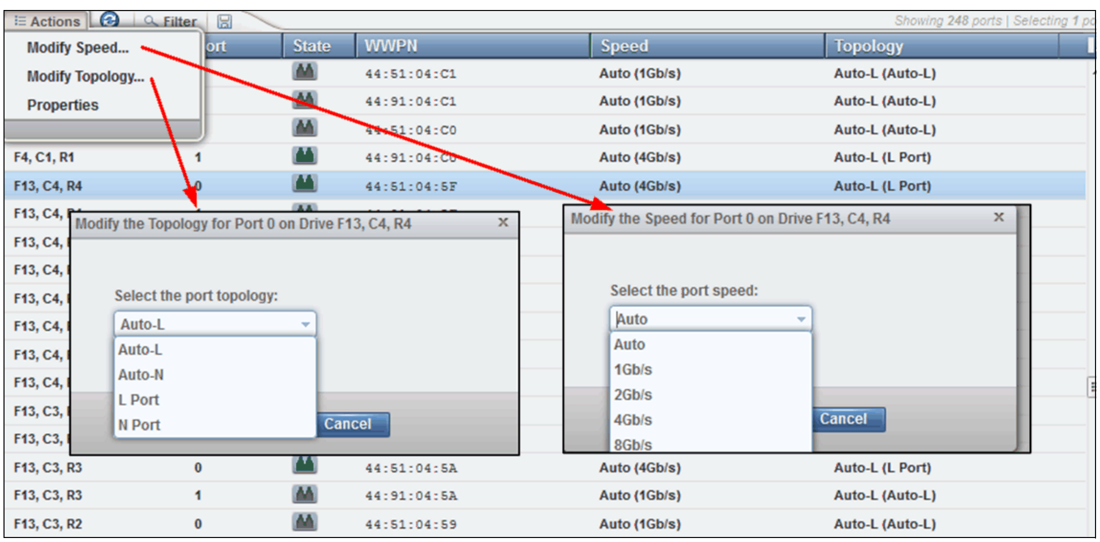

Figure 4-116 Modifying the topology and speed in the Ports view

The following setting values are valid for the topology:

- Auto-L (The drive tries Loop port (L) mode first.)
- Auto-N (The drive tries Node port (N) mode first.)
- L Port (NL)
- N Port

Note: The recommended setting for topology if you are connecting to an FC switch is to use either the Auto-N or N port setting.

The following setting values are valid for the port link speed:

- Auto
- 1 Gbps
- 2 Gbps
- 4 Gbps
- 8 Gbps
- 16 Gbps for TS1160 drive only

From the Actions menu, you can display the properties of a port, as shown in Figure 4-117.

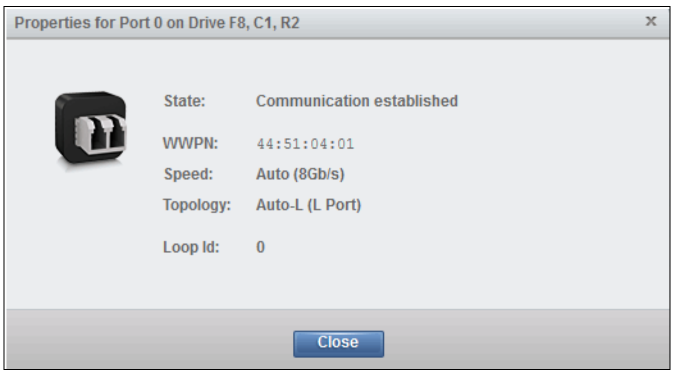

Figure 4-117 Port properties

Exporting drive connectivity report

Administrators can obtain a snapshot of their drive connections by way of the Fibre Ports page or the Ethernet Ports page in a csv file. The common use cases for the new drive connectivity report is for troubleshooting the connections between the tape library and hosts.

Administrators can examine the `.csv` file and find which host is holding reservations along with capturing when the last I/O was performed. Figure 4-118 shows how to export the drive connectivity report from the Fibre Ports page. The SAS and Ethernet Ports pages also download the drive connectivity report in the same location.

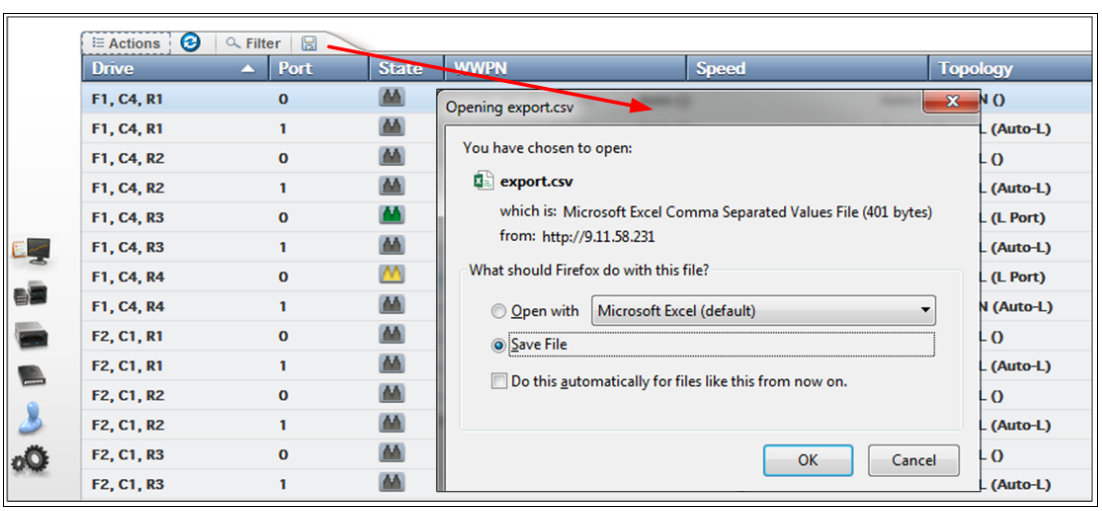

Figure 4-118 Exporting drive connectivity report

Figure 4-119 shows an example of what a drive connectivity report `.csv` file might look like. The data in the file varies per host/tape-library system.

1	Drive S/N	Drive Location	Last Acce:	Drive Port	Transport ID	Reservation	Host Nam	Device Spe	Operating	Operatin
2	00000000A243	F2,C1,R1	09/21/201	2	100090fa559239	None	""			
3	0000078PG255	F2,C1,R2	09/29/201	1	210024ff4fb154	Persist Reserve - Exclusive Access	""			
4	0000078PG251	F2,C1,R3	09/25/201	1	210024ff3d883c	None	""			
5	0000078PG24F	F2,C1,R4	09/29/201	1	210024ff3d898d	Persist Reserve - Exclusive Access	""			
6	0000078PG257	F2,C2,R3	09/29/201	2	210024ff4fb154	Persist Reserve - Exclusive Access	""			
7	0000078DFFD8	F2,C2,R4		1	100090fa559239	None	""			
8	0000078PG144	F3,C1,R2		1	100090fa559239	None	""			
9	0000078PG143	F3,C1,R3		1	100090fa559239	None	""			
10	0000078PG250	F3,C2,R3	09/25/201	1	210024ff3d883c	None	""			
11	0000013B002F	F3,C3,R3	10/01/201	2	100090fa8fd91	None	""			
12	0000013B005E	F3,C3,R4	10/01/201	1	100090fa8fd91	None	""			
13	0000078DFFF5	F3,C4,R1	09/21/201	1	iqn.1994-05.com.redhat:c6165990c8ad,i,0x00023D000002	None	""			
14	0000078DFFE8	F3,C4,R2		1	iqn.1994-05.com.redhat:c6165990c8ad,i,0x00023D000003	None	""			
15	0000078DFFF1	F3,C4,R3		1	iqn.1994-05.com.redhat:c6165990c8ad,i,0x00023D000001	None	""			
16	0000078DFFF0	F3,C4,R4		1	iqn.1994-05.com.redhat:c6165990c8ad,i,0x00023D000004	None	""			
17	0000013B0050	F4,C1,R1	10/01/201	1	210024ff2535f8	None	""			
18	0000013B0043	F4,C1,R2		1	210024ff2535f8	None	""			
19	0000078DBB78	F4,C1,R3	10/01/201	1	100090fa8f11d8	None	""			
20	0000078PG249	F4,C1,R4		1	100090fa559239	None	""			
21	0000078PG249	F4,C1,R4	09/29/201	1	210024ff4fb154	Persist Reserve - Exclusive Access	""			
22	0000078DFFD4	F4,C2,R1		1	100090fa559239	None	""			
23	YD1068015104	F4,C2,R2	09/24/201	2	100090fa8f11d8	None	""			

Figure 4-119 Drive connectivity report

4.6.4 SAS Ports

Use the SAS Ports option in the Drives menu (see Figure 4-120) to view all of the TS1160 model 60S SAS port information for the drives that are in the library. The window shows a summary of all SAS drives, such as state, address, and connection speed. It *cannot* be used to modify any settings.

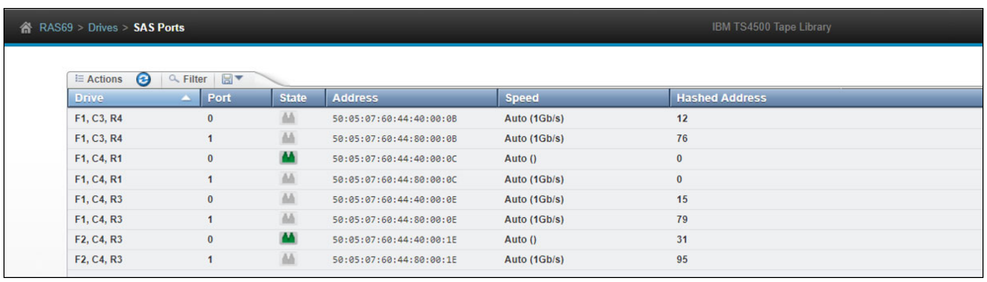

Drive	Port	State	Address	Speed	Hashed Address
F1, C3, R4	0		50:05:07:60:44:40:00:0B	Auto (1Gb/s)	12
F1, C3, R4	1		50:05:07:60:44:80:00:0B	Auto (1Gb/s)	76
F1, C4, R1	0		50:05:07:60:44:40:00:0C	Auto ()	0
F1, C4, R1	1		50:05:07:60:44:80:00:0C	Auto (1Gb/s)	0
F1, C4, R3	0		50:05:07:60:44:40:00:0E	Auto (1Gb/s)	15
F1, C4, R3	1		50:05:07:60:44:80:00:0E	Auto (1Gb/s)	79
F2, C4, R3	0		50:05:07:60:44:40:00:1E	Auto ()	31
F2, C4, R3	1		50:05:07:60:44:80:00:1E	Auto (1Gb/s)	95

Figure 4-120 SAS ports

For more information about exporting the drive connectivity from the SAS Ports page, see “Exporting drive connectivity report” on page 273.

The port State column is sorted by state color (green, yellow, or gray). The gray port status icon is shown in Figure 4-121.

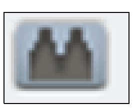

Figure 4-121 Gray port status icon

The gray port status icon indicates that the SAS port is not initialized.

If no light is detected at the drive’s SAS port, one of the following conditions is true:

- An SAS cable is unplugged at the drive or at the nearest SAS device.
- An SAS cable is broken between the drive and the nearest SAS.
- The drive or the nearest SAS device is powered off or has a hardware problem.

The green port status icon is shown in Figure 4-122.

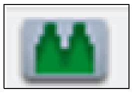

Figure 4-122 Green port status icon

The green port status icon indicates that the SAS port communication is established. The summary display shows information about all of the SAS drives, such as state, address, speed and Hashed address.

From the Actions menu, you can display the properties of a port, as shown in Figure 4-123.

Figure 4-123 SAS port properties

4.6.5 Ethernet Ports

Use the Ethernet Ports option in the Drives menu (see Figure 4-124) to view all of the TS1160 model 60E and TS1155 model 55E Ethernet port information for the drives that are in the library. The window shows a summary of all Ethernet drives, such as state, IP address, iSCSI details, MAC address, and connection speed. It cannot modify any settings. For more information about exporting the drive connectivity from the Ethernet Ports page, see “Exporting drive connectivity report” on page 273.

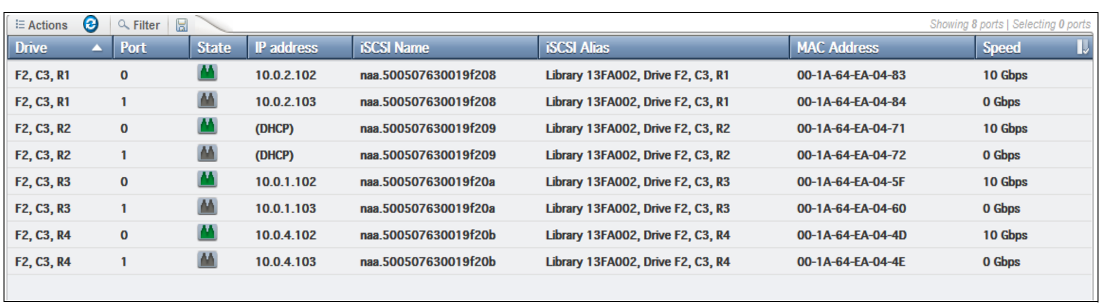

Actions Filter Showing 8 ports | Selecting 0 ports

Drive	Port	State	IP address	iSCSI Name	iSCSI Alias	MAC Address	Speed
F2, C3, R1	0		10.0.2.102	naa.500507630019f208	Library 13FA002, Drive F2, C3, R1	00-1A-64-EA-04-83	10 Gbps
F2, C3, R1	1		10.0.2.103	naa.500507630019f208	Library 13FA002, Drive F2, C3, R1	00-1A-64-EA-04-84	0 Gbps
F2, C3, R2	0		(DHCP)	naa.500507630019f209	Library 13FA002, Drive F2, C3, R2	00-1A-64-EA-04-71	10 Gbps
F2, C3, R2	1		(DHCP)	naa.500507630019f209	Library 13FA002, Drive F2, C3, R2	00-1A-64-EA-04-72	0 Gbps
F2, C3, R3	0		10.0.1.102	naa.500507630019f20a	Library 13FA002, Drive F2, C3, R3	00-1A-64-EA-04-5F	10 Gbps
F2, C3, R3	1		10.0.1.103	naa.500507630019f20a	Library 13FA002, Drive F2, C3, R3	00-1A-64-EA-04-60	0 Gbps
F2, C3, R4	0		10.0.4.102	naa.500507630019f20b	Library 13FA002, Drive F2, C3, R4	00-1A-64-EA-04-4D	10 Gbps
F2, C3, R4	1		10.0.4.103	naa.500507630019f20b	Library 13FA002, Drive F2, C3, R4	00-1A-64-EA-04-4E	0 Gbps

Figure 4-124 Ethernet ports

The port State column is sorted by state color (green, yellow, or gray). The gray port status icon is shown in Figure 4-125.

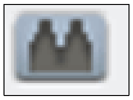

Figure 4-125 Gray port status icon

The gray port status icon indicates that the Ethernet port is not initialized.

If no light is detected at the drive's Ethernet port, one of the following conditions is true:

- A Ethernet cable is unplugged at the drive or at the nearest Ethernet device (host bus adapter, switch, or hub).
- A Ethernet cable is broken between the drive and the nearest Ethernet device (host bus adapter, switch, or hub).
- The drive or the nearest Ethernet device (host bus adapter, switch, or hub) is powered off or has a hardware problem.

The green port status icon is shown in Figure 4-126.

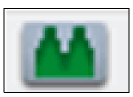

Figure 4-126 Green port status icon

The green port status icon indicates that the Ethernet port initialization is complete.The summary display shows information about all of the Ethernet drives, such as state, IP address, iSCSI details, MAC address and speed of connection.

From the Actions menu, you can display the properties of a port, as shown in Figure 4-127.

Figure 4-127 Ethernet port properties

4.7 Library menu

Use the TS4500 management GUI page under the Library icon to view or manage the logical libraries, as shown in Figure 4-128.

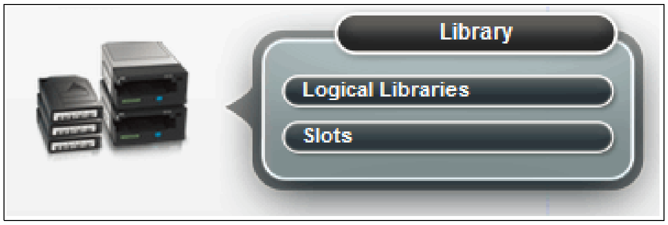

Figure 4-128 Library menu

Note: The Slots option is only available if the user logged in by using the service user ID or if the user has service-level access.

4.7.1 Advanced Library Management System

Advanced Library Management System (ALMS) is an extension of the IBM patented Multi-Path Architecture. With ALMS, the TS4500 tape library can virtualize the locations of cartridges (called *SCSI element addresses*) while it maintains the native storage area network (SAN) attachment for the tape drives. ALMS enables logical libraries to consist of unique drives and ranges of volume serial (VOLSER) numbers, instead of fixed locations.

The TS4500 tape library includes ALMS, so you can immediately start to assign tape drives to any logical library by using the TS4500 management GUI. Logical libraries can also be added, deleted, or easily changed without disruption. Storage capacity can be changed without affecting host applications.

ALMS offers the dynamic management of cartridges, cartridge storage slots, tape drives, and logical libraries. It enables the TS4500 tape library to achieve unprecedented levels of integration for functionality through dynamic partitioning, storage slot pooling, and flexible drive assignment.

ALMS eliminates downtime when you add capacity on demand (CoD) or high-density capacity on demand (HD CoD) storage, add or remove logical libraries, or change logical library storage allocation. ALMS also reduces downtime when you add expansion frames, add or remove tape drives, or change the logical drive allocation.

ALMS provides the following capabilities:

- Dynamic partitioning (storage slot pooling and flexible drive assignment)
- Transparent capability to add or remove storage capacity to any host application
- Capability to configure drives or storage capacity without taking the library offline
- Virtual I/O slots (always enabled) to automatically manage the movement of cartridges between I/O slots and storage slots

The TS4500 tape library complies with the SCSI Medium Changer standard. ALMS is not apparent to the application. The SCSI Medium Changer can be thought of as a “location-centric” interface.

The application that is controlling a SCSI Medium Changer device specifies a source and destination location for each request to move a cartridge. The traditional SCSI library does not have control of the cartridge locations. Instead, the SCSI library acts on behalf of the server.

Functional description

The ALMS features are described in this section.

Storage slot virtualization

The host view of a cartridge location is known as the *SCSI storage element address*. Without ALMS, the storage element address maps directly to a specific storage slot after the library is configured. With ALMS enabled, a specific storage element address is no longer associated with a specific storage slot. Instead, storage slots are virtualized by dynamically associating them with element addresses, as required.

An element address is associated with a storage slot that is selected by the library, as cartridges are moved and inventoried. If a storage element is empty because of a move, that source element address becomes unsolicited. The association of storage element addresses occurs in a way that is transparent to the application software.

You can select the number of storage element addresses for a logical library (as reported to the host application software) by changing the Maximum Number of Cartridges setting for that logical library with the TS4500 management GUI. For each logical library, the default value for this setting is the number of addressable storage slots that is installed in the library for that cartridge type when ALMS is first enabled or after ALMS is enabled when the logical library is created.

Note: You can change the Maximum Number of Cartridges setting for each logical library, but the value must always be greater than or equal to the number of cartridges that is assigned to that logical library. The Maximum Number of Cartridges can be set to a value that is higher than the number of addressable storage slots that is installed at the time to allow future library capacity expansion to not be apparent to the host application software. However, application performance might degrade slightly because of the greater number of addresses. Be careful to not exceed the license limitations of the host application software.

Drive assignment

By using the ALMS flexible drive assignment capability, any drive in any position within any frame can be assigned to any logical library without creating any gaps in drive addresses. Drive (data transfer) element addresses are still mapped to specific drive locations when the drive is assigned. Any drive location can be assigned to any logical library by using the TS4500 management GUI. Each drive that is added to a logical library is assigned to the lowest available element address, regardless of the drive location.

Note: The intermix of models is supported within the same logical library, but only for drives of the same type. LTO drives cannot be mixed with 3592 drives, and 3592 drives cannot be mixed with LTO drives.

ALMS is enabled by default on the TS4500, and the data terminal equipment (DTE) addresses for any newly installed and assigned drives do not depend on the positions of the drives. The DTE address for any newly installed or assigned drive is determined by the sequence in which the drive is assigned to each logical library. With ALMS, drives are assigned to logical libraries by using the Drives by Logical Library page to assign a tape drive to a logical library on the Management Interface (MI).

With this interface, the DTE address for the first drive that is assigned to a new logical library is 256 (x'101'), as shown in Figure 4-129. The DTE address for any other drive that is assigned to a logical library is based on the next available DTE address in that particular logical library. The next available DTE address is the lowest-available DTE address after the starting DTE address. This address fills any gaps that are created when drives are unassigned and removed from a logical library.

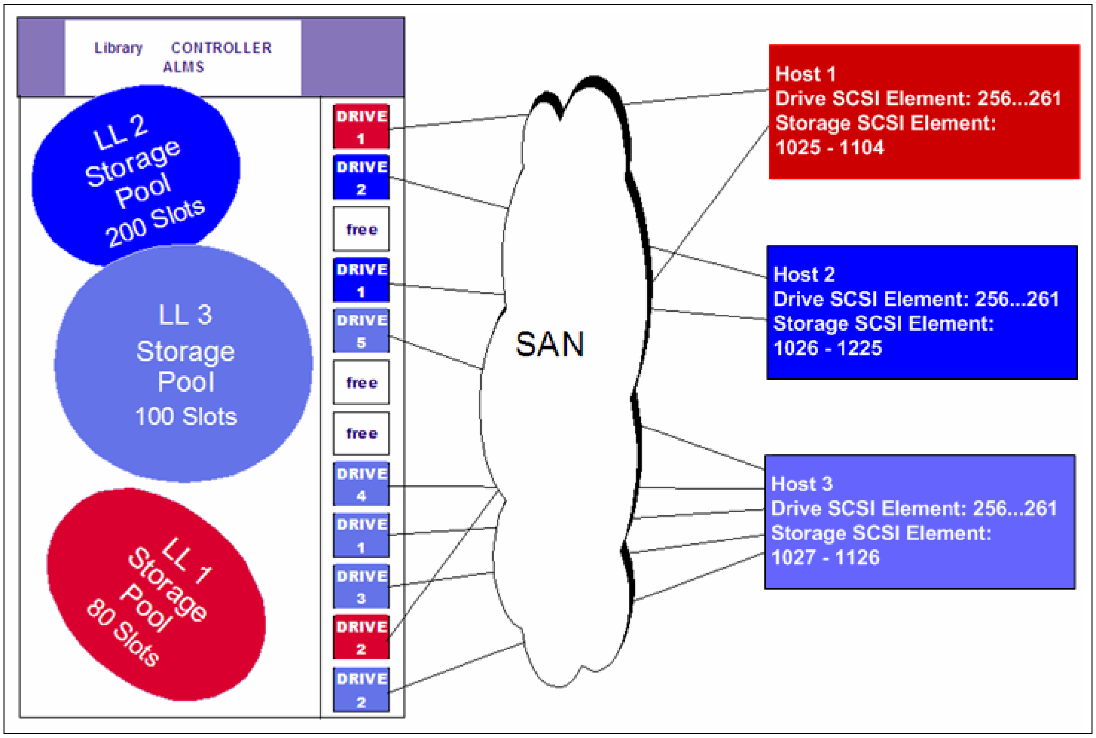

Figure 4-129 TS4500 tape library with ALMS

When a drive is unassigned from a logical library by using the web interface, only that DTE address is made available for future usage. No other DTE addresses are affected.

Storage slot pooling

ALMS allows logical libraries to be added or deleted without disrupting operations. All storage slots are pooled (available on a first-come, first-served basis) to each logical library based on cartridge insert operations. They are a shared resource so that changes to the capacity allocation for each logical library can occur without downtime or administrator involvement.

The minimum logical library has a name only and it can be thought of as a file folder that has no contents. Drives are assigned to the logical library from the Drives page of the TS4500 management GUI. Cartridges are assigned to the logical library based on their VOLSERs and by using one of the following methods, which are listed in priority order:

- VOLSER ranges
- Software application moves from the I/O station (based on the source application that issued the command)
- Manual assignment by an operator by using the TS4500 management GUI

The primary and backup copies of the VOLSER assignment and physical location of cartridges are stored in *nonvolatile random access memory* (NVRAM).

Assigning VOLSER ranges

You can assign specific cartridges to each logical library based on their VOLSERs. With the TS3500 tape library, this function was known as the *Cartridge Assignment Policy*.

All cartridge assignments are displayed on the VOLSER range bar. As you create logical libraries and assign VOLSER ranges, any cartridges that are not part of those ranges remain unassigned. View or assign unassigned cartridges on the Cartridges page.

When you insert a cartridge into the library and its VOLSER is within a range that is assigned to a certain logical library, the cartridge is assigned to that logical library. The cartridge must be of the same media type as that logical library. For example, if you create `Logical Library 1` for VOLSERs that range from `ABC000` to `ABC999` (a library of LTO drives) and then you insert a cartridge with VOLSER `ABC123`, the library recognizes that VOLSER as belonging to the range and assigns it to `Logical Library 1`.

If you insert a cartridge that is outside of any VOLSER ranges, it is available to import into any logical library of the same media type. The assignment is then determined by the first application to import the cartridge.

You can view all of the VOLSER ranges if you select **Actions** → **View VOLSER ranges**. You can also view all VOLSER ranges that are assigned to each logical library on the VOLSER Ranges by Logical Library page.

Virtual I/O

With ALMS, virtual I/O slots enhance the import and export capabilities of the library. With the TS4500 tape library, virtual I/O slots are always enabled so that the library automatically queues all cartridge moves between the I/O station and the storage slots. This capability makes the process of adding and removing cartridges easier and faster.

The TS4500 tape library has I/O stations with cartridge magazines so that you can import and export up to 144 cartridges at any time. The I/O slots are also known as *import/export elements* (IEEs). Virtual I/O slots increase the quantity of available I/O slots by allowing storage slots to appear to the host as I/O slots. These storage slots are also called *virtual import/export elements* (VIEEs).

With virtual I/O slots, the library automatically moves cartridges from the I/O stations to the storage slots, enhancing import and export performance, while also decoupling physical cartridge movement from the application, increasing operator efficiency.

With virtual I/O slots, the library has various mechanisms for selecting the best storage slot location for each inserted cartridge, and the best I/O station for each ejected cartridge. These mechanisms vary depending on the configuration of your library.

The VIEE temporarily takes on the attributes of an IEE until a host moves the cartridge into a *storage element* (StE). When the host move occurs, if the cartridge is in a storage slot, no physical move is needed and the element changes from a VIEE to an StE. Similarly, when a host exports a cartridge from an StE, the physical storage slot is reported as a VIEE without moving the cartridge to the I/O station. The library monitors when free space is available in the I/O station and moves exported cartridges when it is convenient for the library.

If a cartridge cannot be assigned, this situation is reported as `Assignment Pending`. This situation can occur if the assigned logical library does not have any available VIEE slots, or if all of the logical libraries do not have a common VIEE to share. To resolve this situation, free VIEE addresses so that they are available in all libraries or make a specific assignment of this cartridge to a logical library.

Library sharing

The TS4500 tape library can be configured into one or more logical libraries that can be shared by multiple applications.

The guided setup for the TS4500 tape library allows a single application to operate the library. However, often it is advantageous to be able to share a single library between heterogeneous or homogeneous applications. Certain applications (and certain servers) do not allow sharing a library between systems.

With the TS4500 tape library, you can create configurations that enable the library to process commands from multiple *heterogeneous* applications, such as an IBM System p application and a Windows application. With the TS4500 tape library, you can also create configurations that enable the library to process commands from multiple *homogeneous* applications, for example, the same application that is run by several System p servers.

Use the TS4500 management GUI to perform the following actions.

Separate libraries to separate applications

Configure the library so that it is partitioned into separate logical libraries that independently communicate with separate applications through separate control paths. This configuration (as shown in Figure 4-130) requires no special capabilities from the server or application.

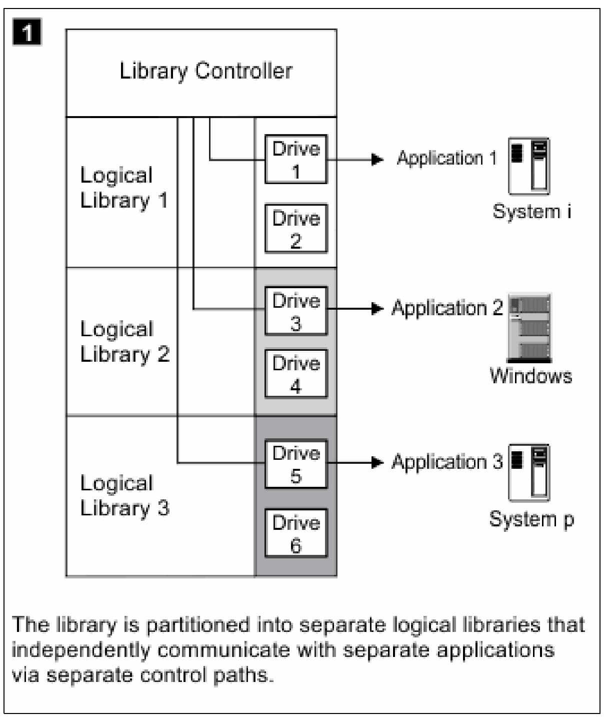

Figure 4-130 Separate libraries to separate applications

Configure any single logical library (including the entire physical library) so that it is shared by two or more servers that are running the same application. Depending on the capabilities of the server and application, several ways exist to set up this type of configuration. The following methods are typical examples:

- One server to a single control path
- All servers to a single control path
- Multiple servers through multiple control paths

One server to a single control path

Configure one server (mainframe host) to communicate with the library through a single control path. All other servers send requests to that server through a network (as shown in Figure 4-131). This configuration is used by IBM Spectrum Protect.

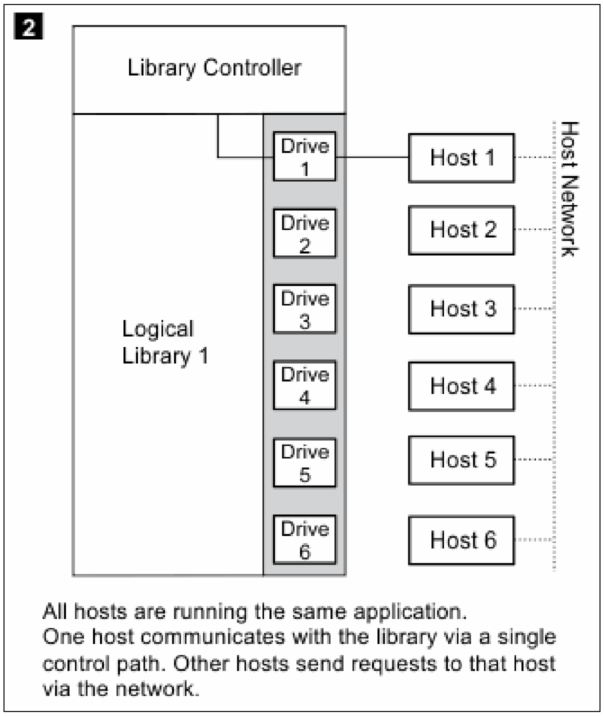

Figure 4-131 Separate libraries to separate applications

All servers to a single control path

Configure all of the servers to communicate with the library through a single, common control path (as shown in Figure 4-132 on page 283). This configuration is used in high-availability environments, such as IBM High Availability Cluster Multi-Processing (IBM HACMP) from IBM, and Systems Management Server and Clustered Server Environments from Microsoft. Multi-initiator configurations are supported only by certain adapters and independent software vendors (ISVs). For more information, contact your vendor.

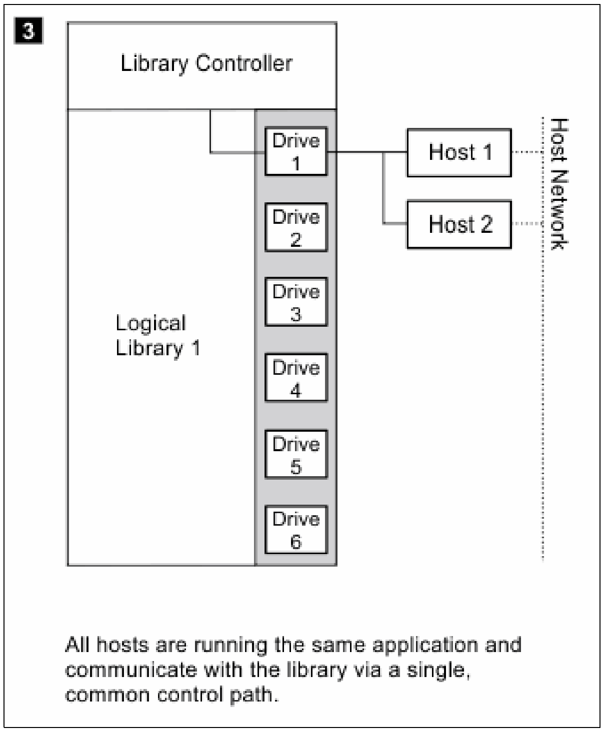

Figure 4-132 Separate libraries to separate applications

Multiple servers through multiple control paths

Configure multiple logical libraries to communicate with multiple servers through multiple control paths. This configuration (see Figure 4-133) requires that you add multiple logical libraries, and it is used by separate Backup, Recovery, and Media Services.

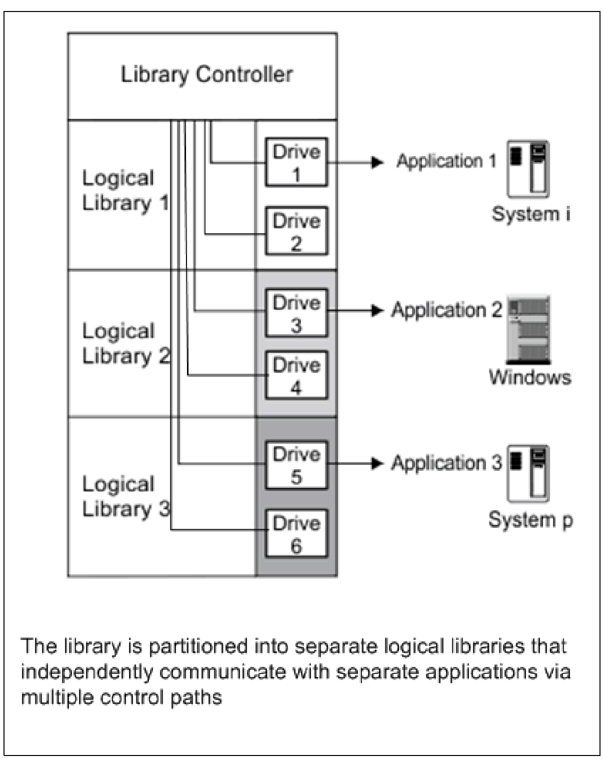

Figure 4-133 Separate libraries to separate applications

Your library configuration is not limited to these examples. Many configurations are possible and you can design them according to your business needs.

4.7.2 Creating Logical Library window

Use this window to manage and reconfigure logical libraries, including modifying the method of encryption that is used on the logical library. Figure 4-134 shows a list of configured libraries and the window that is used to set the columns of information to display.

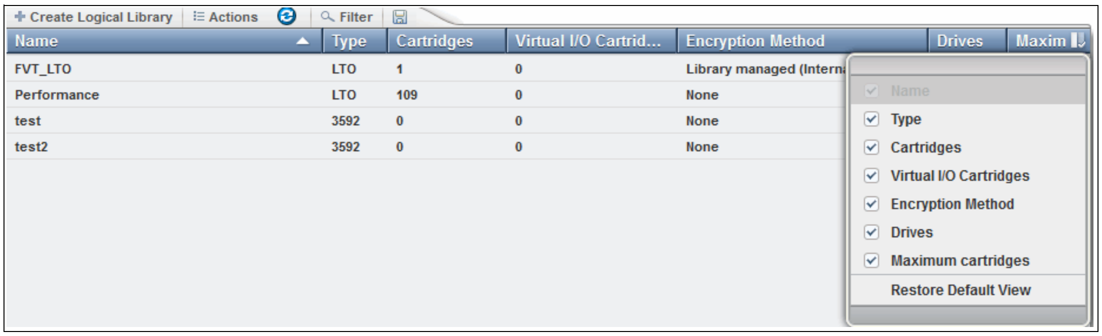

Figure 4-134 Logical Library main window

Using multiple logical libraries

The TS4500 tape library supports multiple libraries to share the physical library between applications or to support mixed drive types for any application.

You can create multiple logical libraries by partitioning the library's tape drives and tape cartridges into two or more logical libraries. Each logical library consists of the following components:

- Tape drives
- Tape cartridges in storage slots
- Tape cartridges in input/output (I/O) slots

Each logical library has its own control path. A *control path* is a logical path into the library through which a server sends standard SCSI Medium Changer commands to control the logical library. For frames that contain LTO tape drives and frames that contain 3592 tape drives, each logical library control path is available to servers through logical unit number (LUN) 1 of the first drive that is defined within that logical library. A *LUN* is a number that is used by a server to identify a drive.

In a TS4500 tape library, a logical library shares empty storage slots, empty I/O slots, and the cartridge accessor on a first-come, first-served basis.

> **Note:** When you name logical libraries, develop an enterprise-wide standard or convention for naming libraries for easier management of the libraries.

With automatic cleaning, cleaning cartridges are shared among logical libraries, so any appropriate cleaning cartridge is used to clean a drive in any configured logical library.

You can create multiple logical libraries by partitioning the physical library's tape drives and tape cartridges into two or more logical libraries. Each logical library can contain only one media type.

Creating logical libraries

Create logical libraries by using presets that have the preferred drive and encryption configurations for their specific use. Select the **Create Logical Library** tab to see a selection of default logical library options, including a custom option, as shown in Figure 4-135.

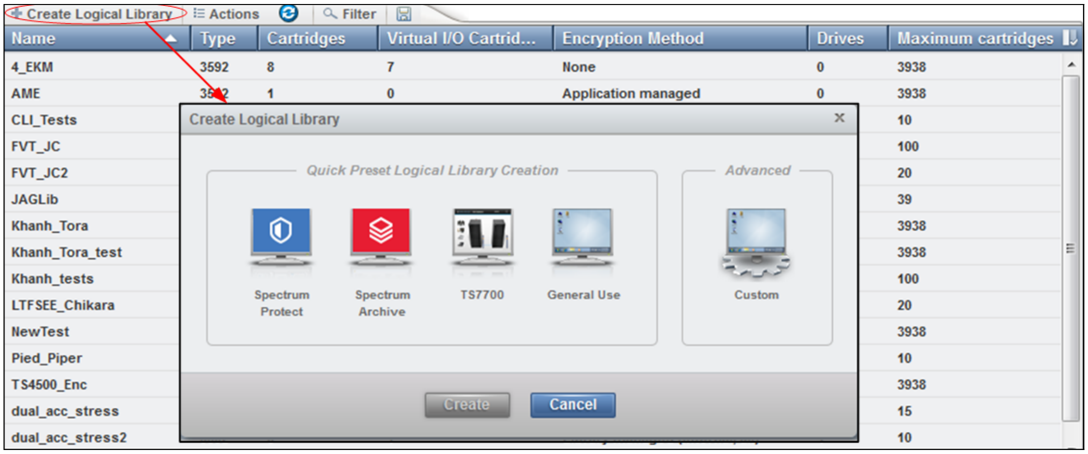

Figure 4-135 Create Logical Library

Spectrum Protect

Use this option if you manage your library with Spectrum Protect.

When you select this option, 2 - 4 of the drives that are assigned to the logical library are designated as control path drives. When possible, two consecutive drives are assigned as control paths in two separate frames. This selection configures all unassigned drives to this logical library.

To configure a logical library by using the Spectrum Protect option, select the **Spectrum Protect** icon, then define a name, select the media drive type (LTO or 3592), and select the encryption method to use. If necessary, enter the VOLSER range to use with this logical library in the Identifier field, and select **Create**, as shown in Figure 4-136.

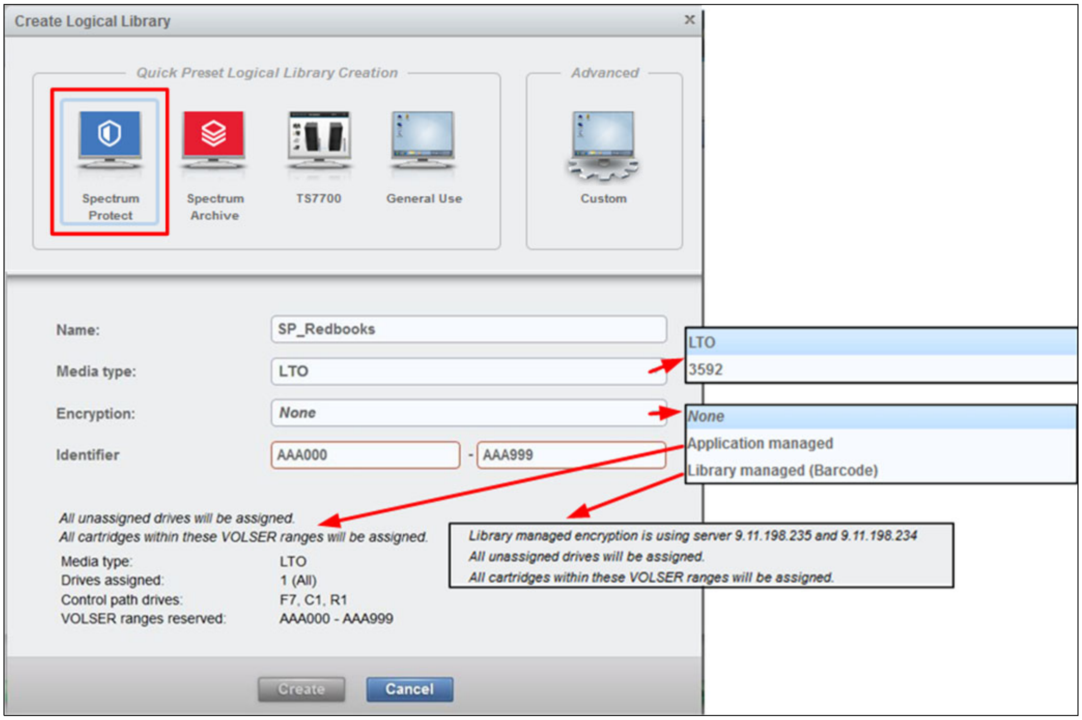

Figure 4-136 Logical library with Spectrum Protect

If library-managed encryption is used, the key server can be set up in the security option that is described in “Encryption key servers” on page 222. Also, the key server can be added by using the Modify Encryption Method option from the Actions menu, as described in “Modify Encryption Method option” on page 293.

To change drive assignments or control paths, see 4.6.1, “Drives” on page 255.

IBM Spectrum Archive

Use this option if you plan to use IBM Spectrum Archive, incorporating the Linear Tape File System (LTFS) format standard for reading, writing, and exchanging descriptive metadata on formatted tape cartridges.

When you select this option, 2 - 4 of the drives that are assigned to the logical library are designated as control path drives. When possible, two consecutive drives are assigned as control paths in two separate frames. This selection configures all unassigned drives to this logical library.

To configure a logical library by using the Spectrum Archive option, select the **Spectrum Archive** icon, then define a name, select the media drive type (LTO or 3592), and select the encryption method to use, if needed. Insert the VOLSER range to use for this logical library in the Identifier field and select **Create**, as shown in Figure 4-137.

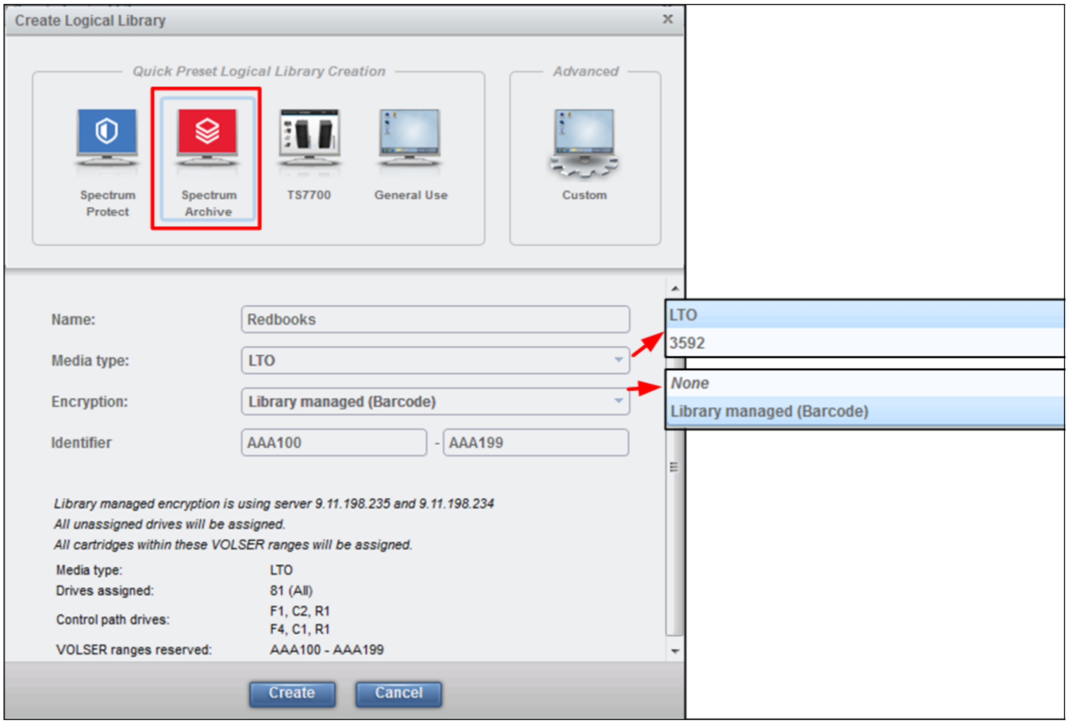

Figure 4-137 Logical library with IBM Spectrum Archive

If library-managed encryption is used, the key server can be set up in the security option that is described in “Encryption key servers” on page 222. The key server can be added by using the **Modify Encryption Method** option from the **Actions** menu, as described in “Modify Encryption Method option” on page 293.

For more information about changing drive assignments or control paths, see 4.6.1, “Drives” on page 255.

TS7700

Use this option to create a logical library for the IBM Virtualization Engine TS7700 by using all unassigned 3592 tape drives. When you select this option, up to four of the drives that are assigned to the logical library are designated as control path drives. When possible, drives are assigned as control paths in two separate frames for redundancy. The TS7700 option is shown in Figure 4-138.

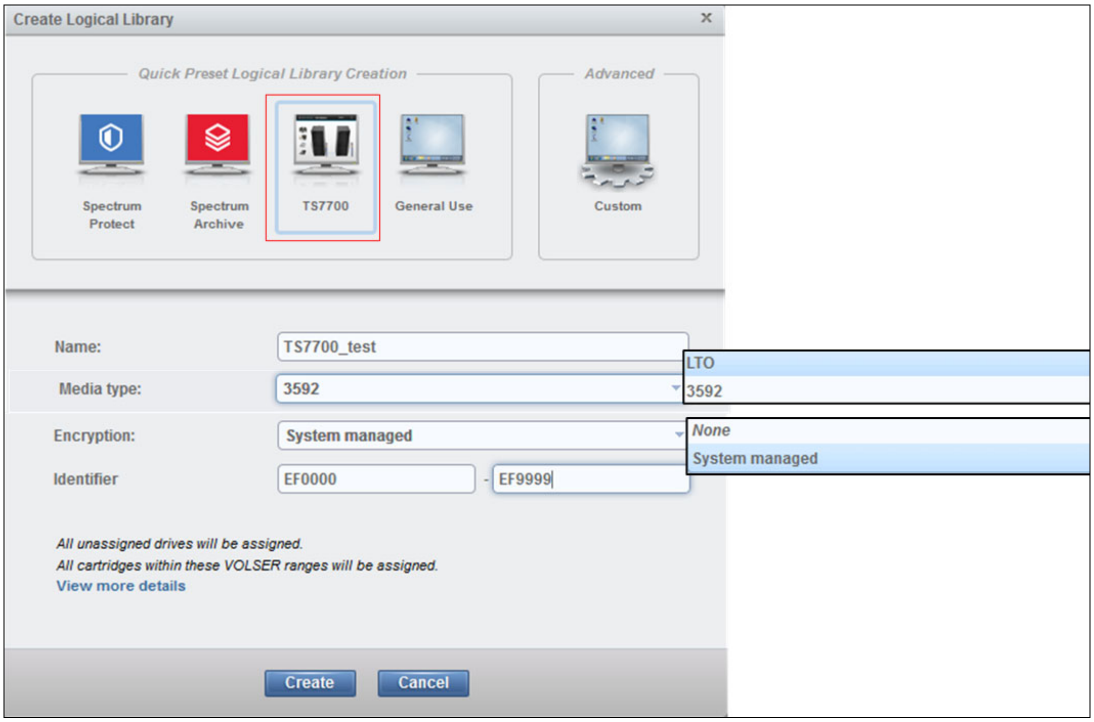

Figure 4-138 Logical Library for TS7700

Note: The TS7700 is only supported by 3592 drive types, and the encryption method must be set to System managed. The TS7700 must run R4.0 or higher microcode.

General Use

Use the General Use option if you have another application to manage your TS4500 tape library and you want to use the default options.

When you select this option, one of the drives that is assigned to the logical library is designated as a control path drive. This selection configures all unassigned drives to this logical library.

To configure a logical library by using the General Use option, select the **General Use** icon, then define a name, select the media drive type (LTO or 3592), and select the encryption method to use, if needed. Insert the VOLSER range to use for this logical library in the Identifier field and select **Create**, as shown in Figure 4-139.

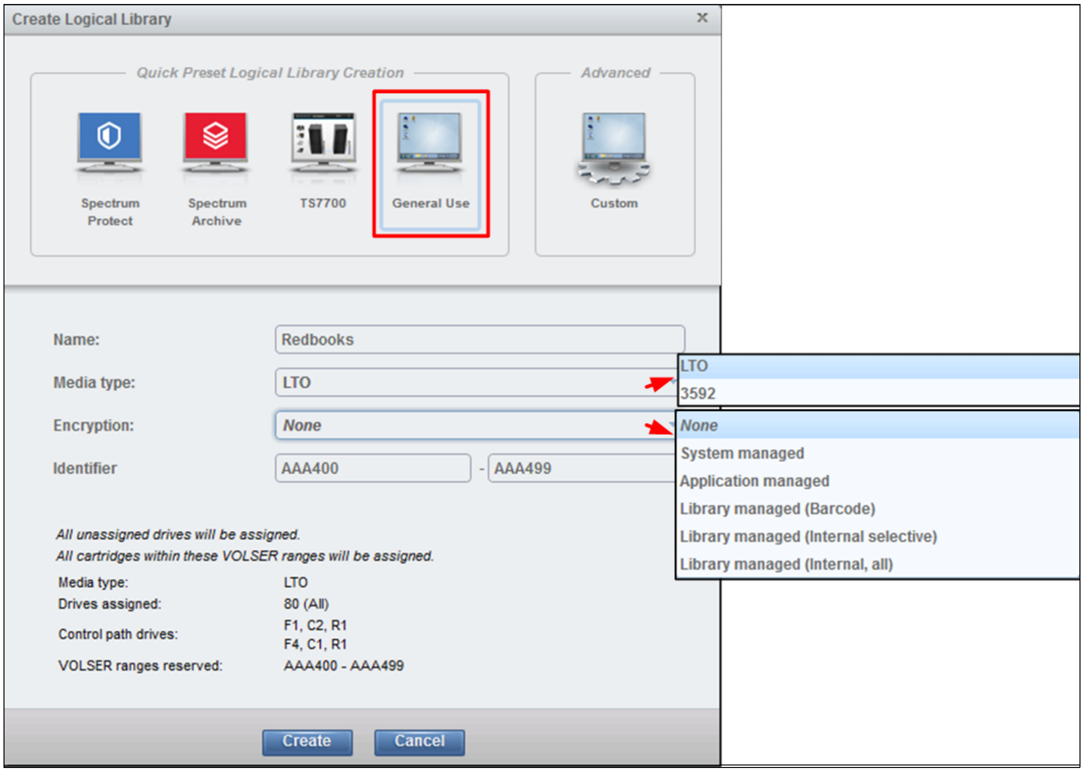

Figure 4-139 Logical library with General Use option

Custom

Use this option for a complete customization when you create a logical library if the presets do not fit your requirements. Use the Custom option to set library information and select the drives, VOLSER range, and maximum number of cartridges. Use this option if you are setting up multiple logical libraries.

Complete the following steps to use the Custom option:

1. Select the number of drives and slots to use for the logical library that you are configuring. Figure 4-140 shows the initial menu after you select **Custom**. Additionally, you can select the drive assignments, VOLSER range, maximum cartridges, and summary.

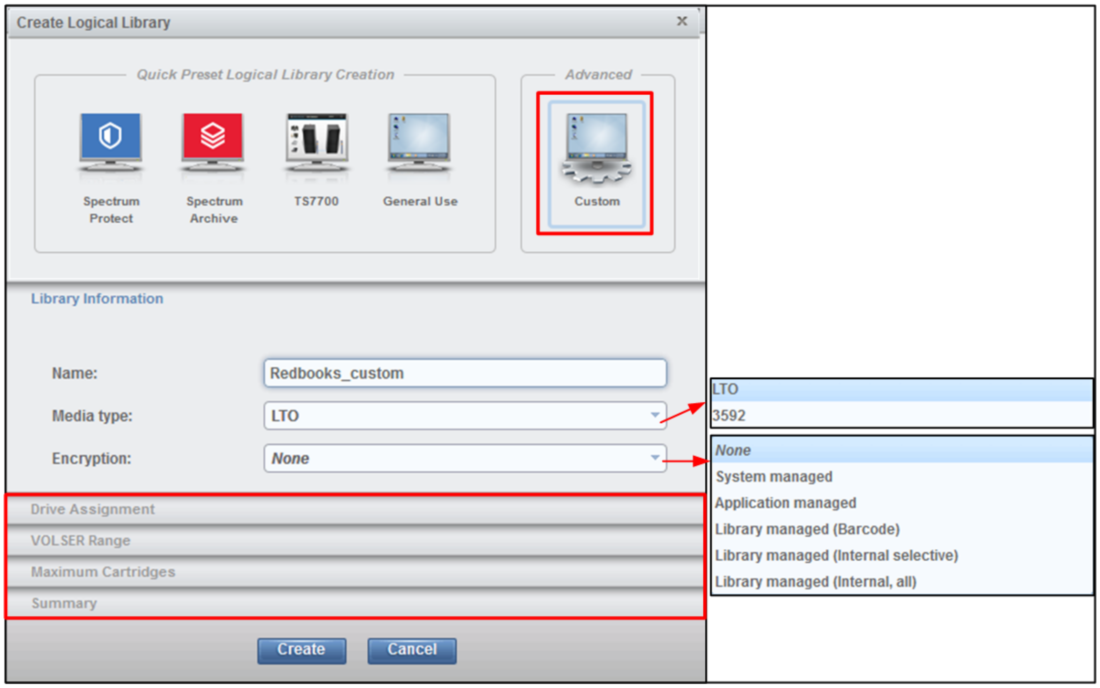

Figure 4-140 Custom logical library

2. Use the first window to define a name, select the media drive type (LTO or 3592), and select the encryption method to use, if required.
3. Use the Drive Assignment page to configure drives for the logical library, as shown in Figure 4-141. You can select all unassigned drives or select only the drives that you need. Only the unassigned drives of the type that was selected (LTO or 3592) are displayed.

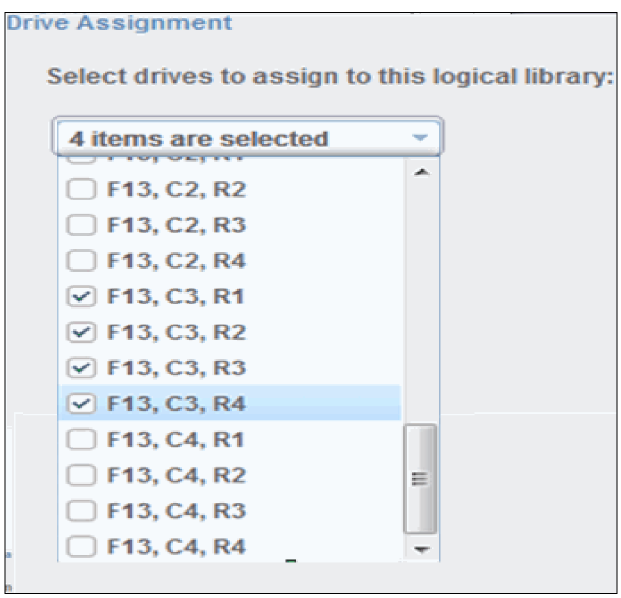

Figure 4-141 Custom drive settings

4. Use the VOLSER Range page to set up a VOLSER's range and select whether these range settings apply to only newly inserted cartridges, as shown in Figure 4-142.

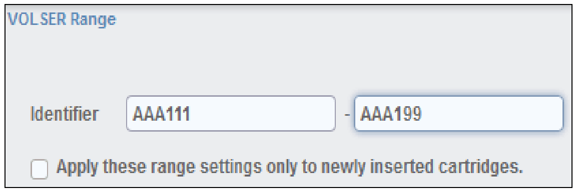

Figure 4-142 Custom VOLSER range

5. Use the Maximum Cartridges page to set the number of cartridge slots that is assigned to this logical library. Also, you can use this page to configure the number of slots by using the system default, as shown in Table 4-143. The system default is the total number of slots that is licensed in the library.

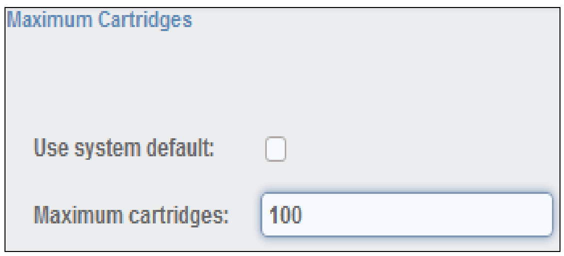

Figure 4-143 Custom maximum cartridges

6. Use the Summary page to display all of your settings, as shown in Figure 4-144.

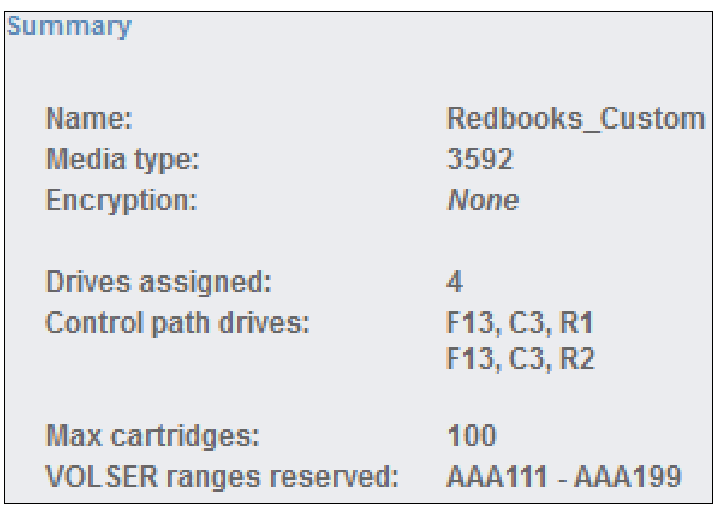

Figure 4-144 Summary

Assigning tape drives to a logical library

Each logical library requires at least one tape drive as a control path drive. Assign more tape drives, as needed, to improve performance. Enable or disable control path drives on the Drives page, as described in "Adding or removing tape drives" on page 256.

Drives are assigned to logical libraries when the logical libraries are created. You can modify the drives that are assigned to a logical library on the Drives by Logical Library page, as described in "Reassigning a drive" on page 257.

Note: A tape drive can be used only by the logical library to which it is assigned. Drives cannot be shared between logical libraries.

Assigning cartridges to a logical library

Cartridges are assigned to logical libraries based on the VOLSER range that is assigned to the logical library when it is created. For more information, see "Assigning VOLSER ranges". You can modify the cartridges that are assigned to a logical library by modifying the VOLSER range as described in "Creating or modifying a VOLSER range" on page 252.

Note: Cleaning cartridges are shared among the logical libraries.

Assigning VOLSER ranges

You can assign specific cartridges to each logical library based on their volume serial (VOLSER) numbers. With the TS3500 tape library, this function was known as the *Cartridge Assignment Policy*.

All cartridge assignments are displayed on the VOLSER range bar. As you create logical libraries and assign VOLSER ranges, any cartridges that are not part of those ranges remain unassigned. View or assign unassigned cartridges on the Cartridges page, as described in 4.5.1, "Cartridges" on page 238.

When you insert a cartridge into the library and its VOLSER is within a range that is assigned to a certain logical library, the cartridge is assigned to that logical library. The cartridge must be the same media type as the media type of that logical library. For example, if you create `Logical Library 1` for VOLSERs that range from `ABC000` to `ABC999` (a library of LTO drives), and then you insert a cartridge with VOLSER `ABC123`, the library recognizes that VOLSER as belonging to the range and assigns it to `Logical Library 1`.

If you insert a cartridge that is outside of any VOLSER ranges, the cartridge is available to import into any logical library of the same media type. The assignment is then determined by the first application to import the cartridge.

You can view all of the VOLSER ranges by selecting **Actions** → **View VOLSER ranges**. You can also view all VOLSER ranges that are assigned to each logical library on the VOLSER Ranges by Logical Library page, as described in 4.5.2, "Cartridges by Logical Library" on page 247.

Create Logical Library Actions menu

After a logical library is configured, you can use the **Actions** menu (see Figure 4-145) to select a configured logical library and modify it.

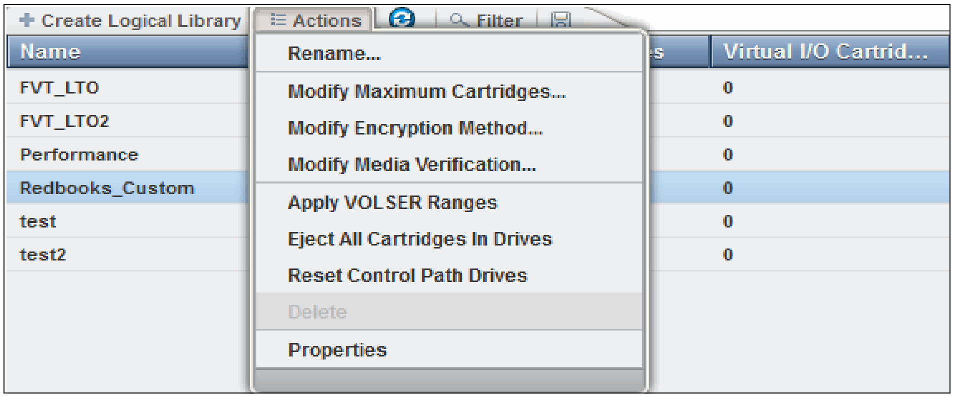

Figure 4-145 Logical Library Actions menu

Rename

Use the Rename option in the Actions menu to rename a logical library.

Select the logical library. Click **Actions** → **Rename**. The Rename window opens and you can enter a new logical library name.

Modify Maximum Cartridges option

Use the Modify Maximum Cartridges option on the Actions menu to modify the maximum number of cartridges that is configured in a logical library. If a logical library ran out of storage capacity or new frames were added, this option can be used to increase or decrease the logical library cartridge capacity.

Select the logical library and click **Actions** → **Modify Maximum Cartridges**. The Modify Maximum Cartridges window opens and you can modify the number of cartridge slots for the logical library. If you added a storage frame, you can also use the Reset to system default option, which automatically enters the full licensed capacity of all slots in the physical library.

Modify Encryption Method option

Use the Modify Encryption Method option on the Actions menu to modify the encryption method and to configure the EKM servers to use on a logical library.

Select the logical library, and click **Actions** → **Modify Encryption Method**. The Modify Encryption Method window opens so that you can set up or remove encryption from the logical library. With library-managed encryption, you can use the Modify Encryption Method page to set the EKM server IP address to use on this logical library. A drop-down list shows all of the configured EKM servers.

You can add up to four servers to the logical library, as shown on Figure 4-146. You can use this page to add a new unlisted EKM server IP address and to test it by clicking **Ping**.

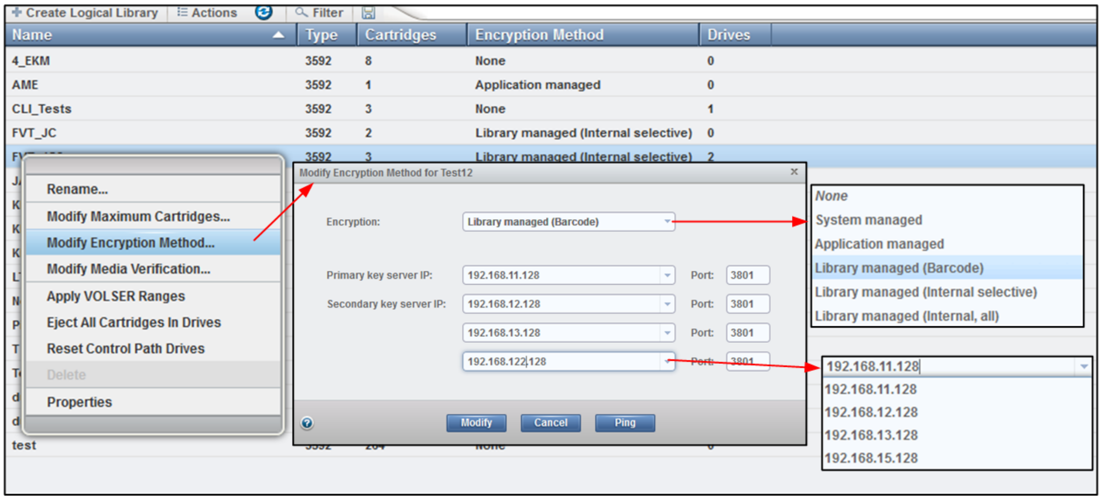

Figure 4-146 Modifying the encryption method

Apply VOLSER Ranges option

Use the Apply VOLSER Ranges option on the Actions menu to reassign all cartridges based on VOLSER ranges. This action can be useful in the following scenarios:

- If cartridges with old policies that do not match the current assignments are in the library
- If a VOLSER range was moved to a different logical library, or the range was resized and the option to reassign cartridges was not selected
- If a cartridge was manually assigned outside of a VOLSER range

Cartridges are reassigned to existing logical libraries. If the maximum number of cartridges is met for a logical library, any other cartridges that were assigned to that logical library are unassigned.

Eject All Cartridges in Drives option

Use the Eject All Cartridges in Drives option in the Actions menu to eject all cartridges from all drives in a logical library. Use this option only if the cartridges cannot be ejected by the application.

Reset Control Path Drives option

Use the Reset Control Path Drives option on the Actions menu to reset all control path drives that are configured on a logical library.

This option removes all paths on the selected logical library to the TS4500 for the time that it takes to reset the control path drives.

Delete option

Use the Delete option on the Actions menu to delete the logical library.

To delete a logical library, you must perform the following tasks:

- Unassign all drives (see “Unassigning a drive” on page 257)
- Delete all VOLSER ranges (see “Deleting a VOLSER range” on page 253)
- Unassign any cartridges that are assigned (see “Assigning or unassigning cartridges” on page 249)

After you complete these actions, select the logical library to delete. Click **Actions** → **Delete** and follow the instructions on the window.

Properties

Click **Actions** → **Properties** to display detailed properties about the selected logical library, as shown in Figure 4-147.

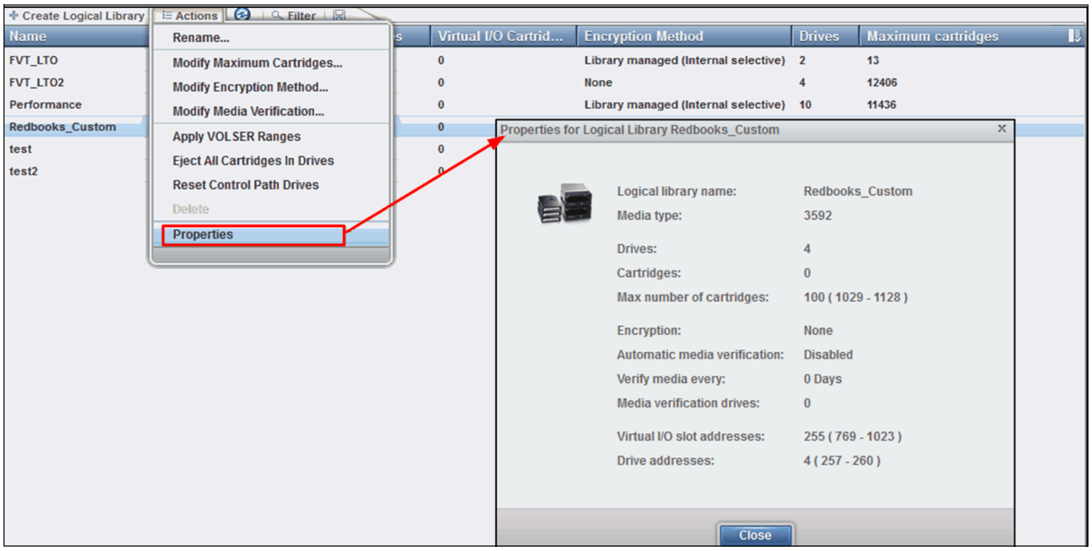

Figure 4-147 Properties

4.7.3 Slots

The Slots option is available for a user with service-level access. Use the Slots option to set storage slots online or offline.

This option is normally used by the IBM service support representative (SSR) for setting slots offline and then online during a service action. This option can also be used to set any slot online and offline for problem determination.

To set slots offline or online, select the frame, row, and slot number from the menu, as shown in Figure 4-148, and then select either the Set Online action or the **Set Offline** action.

Single or multiple slots can be selected and the number that was selected is displayed in the upper-right corner of the menu.

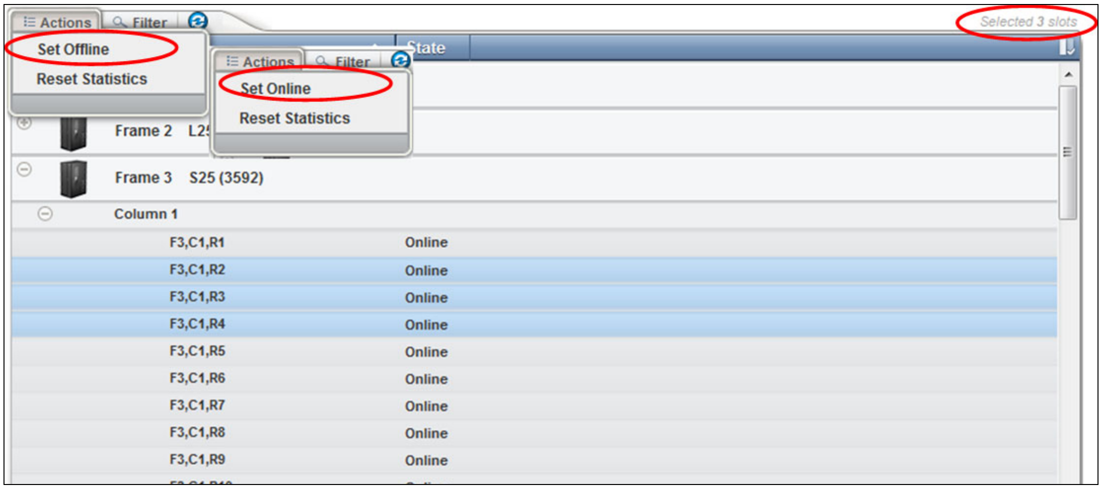

Figure 4-148 Slots

4.8 Monitoring

The TS4500 management GUI pages under the Monitoring icon can help you to monitor the system, events, and tasks, as shown in Figure 4-149.

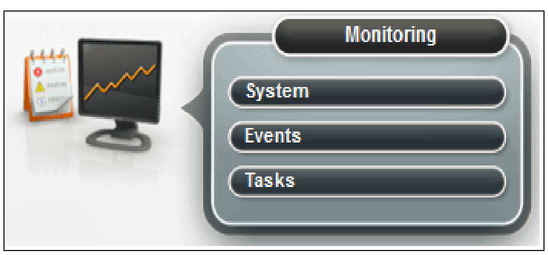

Figure 4-149 Monitoring

4.8.1 System

The System page, as shown in Figure 4-150, displays the configuration (firmware version, Model, SN, and number of frames) and the recent Activity of your TS4500 tape library. Start at this page to identify health and status issues with the library and library components.

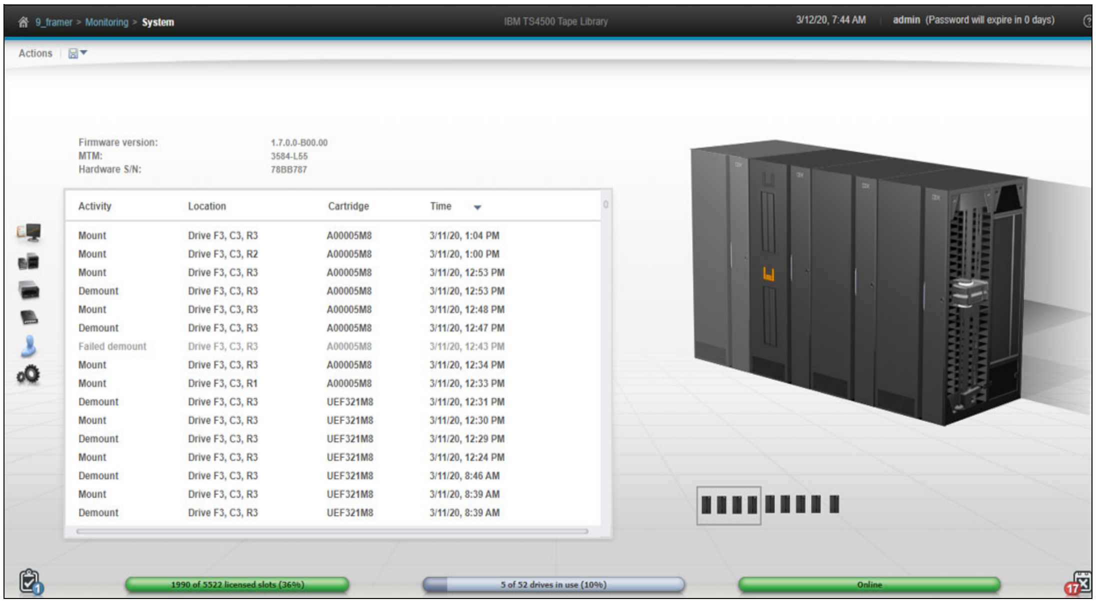

Figure 4-150 System

Viewing Activity Log

The recent Library Activity Type, Location, Cartridge Volser, and Time is shown in Figure 4-151.

Activity	Location	Cartridge	Time
Mount	Drive F3, C3, R1	A00001M8	3/25/20, 9:58 AM
Mount	Drive F3, C3, R2	UEF329M8	3/25/20, 9:57 AM
Demount	Drive F3, C3, R1	UEF329M8	3/25/20, 9:57 AM
Move	Slot F2, C6, R13, T0	UEF328M8	3/25/20, 9:56 AM
Move	Slot F2, C6, R25, T0	UEF327M8	3/25/20, 9:56 AM
Mount	Drive F3, C3, R1	UEF329M8	3/25/20, 9:56 AM
Mount	Drive F2, C1, R4	UEF344M8	3/25/20, 9:55 AM
Mount	Drive F3, C3, R3	A00001M8	3/25/20, 9:54 AM
Mount	Drive F2, C1, R4	A00001M8	3/25/20, 9:53 AM
Mount	Drive F3, C3, R1	UEF344M8	3/25/20, 9:52 AM
Move	Slot F2, C6, R7, T0	A00005M8	3/25/20, 9:51 AM
Mount	Drive F2, C1, R4	UEF344M8	3/25/20, 9:51 AM
Demount	Drive F3, C3, R1	UEF327M8	3/25/20, 9:51 AM
Mount	Drive F3, C3, R1	UEF327M8	3/25/20, 9:49 AM
Demount	Drive F2, C1, R4	UEF329M8	3/25/20, 9:49 AM
Mount	Drive F2, C1, R4	UEF329M8	3/25/20, 9:48 AM
Demount	Drive F3, C3, R1	UEF328M8	3/25/20, 9:47 AM
Move	Slot F2, C6, R7, T0	A00005M8	3/25/20, 9:47 AM

Figure 4-151 Activity Log

Activity types include:

Mount A cartridge was inserted into a drive.

Unmount A cartridge was removed from a drive.

Import A cartridge was inserted into the library by way of I/O slot.

Export A cartridge was removed from the library by way of I/O slot.

Move A cartridge was moved from one slot to another.

Open A library door or I/O station was opened.

Close A library door or I/O station was closed.

Viewing component status

View the status of the library frames, storage slots, tape drives, I/O stations, and the cartridge accessors on the System page.

Hover over the top of a frame, or over an I/O station, to view its state and contents, as shown in Figure 4-152.

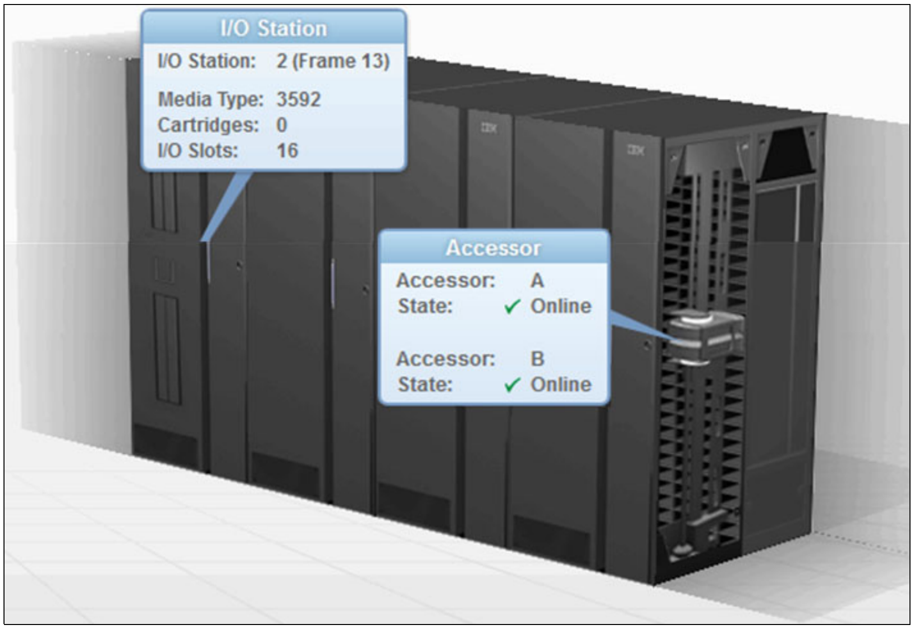

Figure 4-152 Component status

If a problem occurs with any of the components, a warning or error icon is displayed. Hover over the icon to see a message about the affected component that includes the amount of time that passed since the event occurred and a short description of the event. Go to the Events page to view the detailed information about the error or warning.

Library frame

An error or a warning icon at the base of a library frame indicates a problem with the frame door, I/O station (if applicable), or with any of the slots, cartridges, or drives in that frame.

Cartridge accessor

An error or a warning icon on the cartridge accessor indicates a degraded state of operation for the accessor or an issue with any of the accessor components (gripper, scanner, or calibration sensor).

I/O station

Click an I/O station to view the I/O station properties, including the state of the I/O station. The possible states are empty, cartridges in I/O for 60 minutes or more, or a magazine that is missing for 60 minutes or more. Icons on the door of the library frame also indicate the I/O station status. Each icon is described next.

Icon meanings

The System page shows a graphical representation of your library. The icons provide basic status and event information. Hover over error and warning icons to view the length of time since the event was generated and a short description of the event.

Table 4-7 lists each icon.

Table 4-7 Status icons

Icon	Description
	An error occurred with a frame, tape drive, cartridge, or accessor. Hover over the icon to see a list of the most important issues that cause this state. Click any issue to open the Events Page to see more information.
	A warning occurred with a frame, tape drive, cartridge, or accessor. Hover over the icon to see a list of the most important issues that cause this state. Click any issue to open the Events page to see more information.
	The Tasks icon shows whether tasks are running.
	The Alerts icon shows whether alerts exist.
	The I/O station is full.
	The I/O station is full for 60 minutes or longer. This situation might occur with input operations if a shortage of storage slots exists. This situation also might occur with output operations if the I/O station is unattended. View the state of the cartridges, I/O slot (import queued), or I/O slot on the Cartridges page.
	The I/O station is partially full.
	The I/O station is partially full for 60 minutes or longer. This situation might occur with input operations if a shortage of storage slots exists. This situation might occur with output operations if the I/O station is unattended. View the state of the cartridges, I/O slot (import queued), or I/O slot on the Cartridges page.
	The I/O station is empty.
	The cartridge magazine is missing.
	The cartridge magazine is missing for 60 minutes or longer.

Frame selection

The main System status display shows up to four frames at a time. Use the mouse to select the four frames that you want to display from the System status display window, as shown in Figure 4-153.

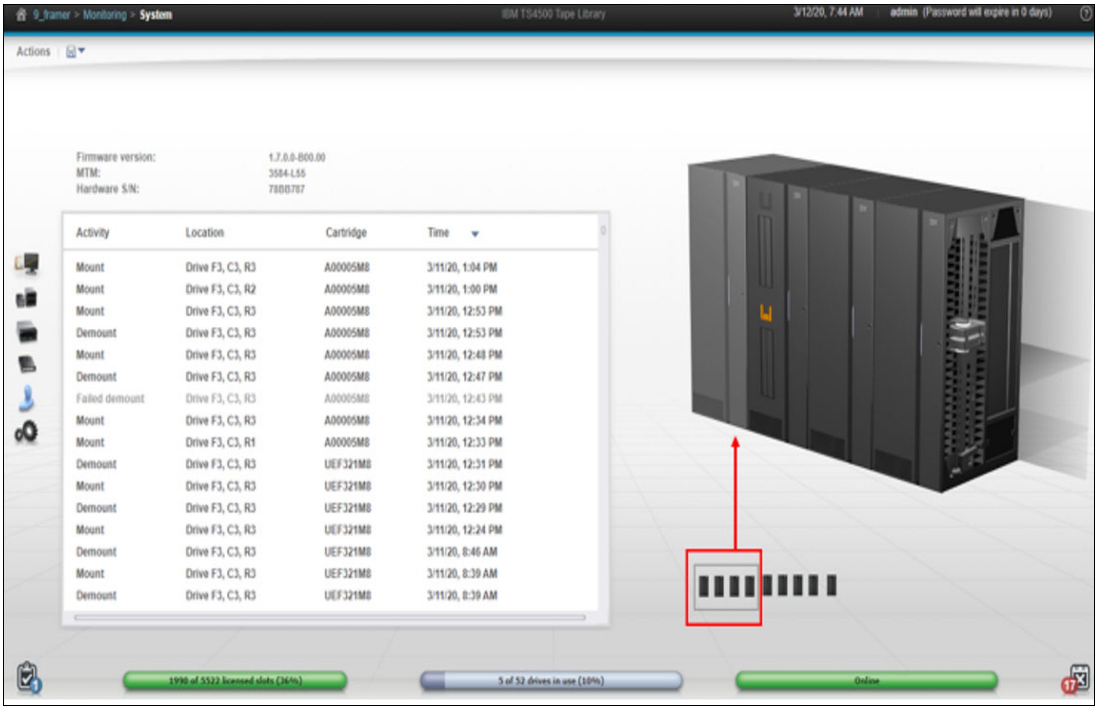

Figure 4-153 Frame selection

Status pods

Status pods at the bottom of the System page show a quick view of the capacity, drive use, and logical library health status, as shown in Figure 4-154.

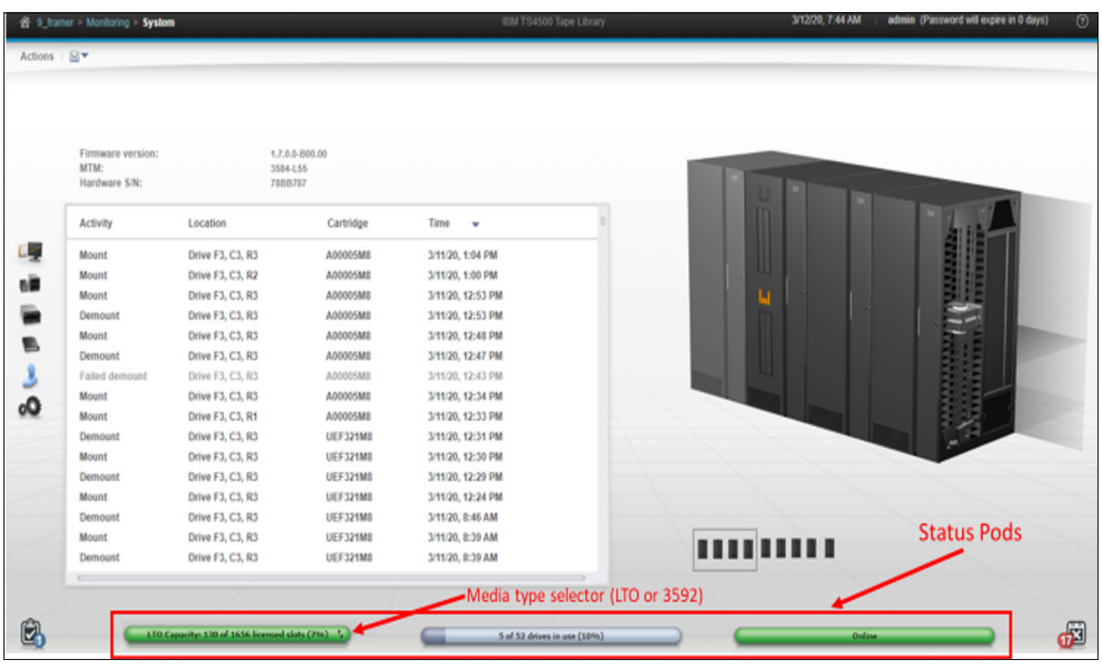

Figure 4-154 Status pods

Physical capacity status pod

The physical capacity pod displays how many licensed slots are filled by cartridges in this tape library. In a mixed media library, you can select to display a summary of either LTO or 3592 slots by using the selector button on this pod.

The pod is green when the number of cartridges in the library is fewer than the number of licensed slots.

The pod is yellow when the physical capacity exceeds the capacity use threshold. If this situation occurs, you can remove cartridges or purchase additional capacity.

The pod is red when the number of cartridges exceeds the number of licensed slots. If this situation occurs, you must remove cartridges or purchase additional capacity.

Drive use status pod

The drive use status pod displays the status of the drives across the library. Click the **status pod** to go to the Drives by Logical Libraries page.

The blue pod fills as the number of drives (in use out of the total number of drives in the library) increases.

If multiple logical libraries exist, a warning icon displays if all of the drives in any logical library are in use.

Health status pod

The color of the health status pod indicates the current state of the library by severity. Click the **status pod** if the pod is yellow or red to go to the correct page so that you can understand and troubleshoot the warning or error. If a library is in more than one state, the higher priority state is shown.

Actions and remote management of the library

Use the Actions menu on the upper-left corner of the page, or right-click a frame or I/O station (see Figure 4-155). The menu of available actions changes depending on whether the entire library, a specific frame, or an I/O station is selected.

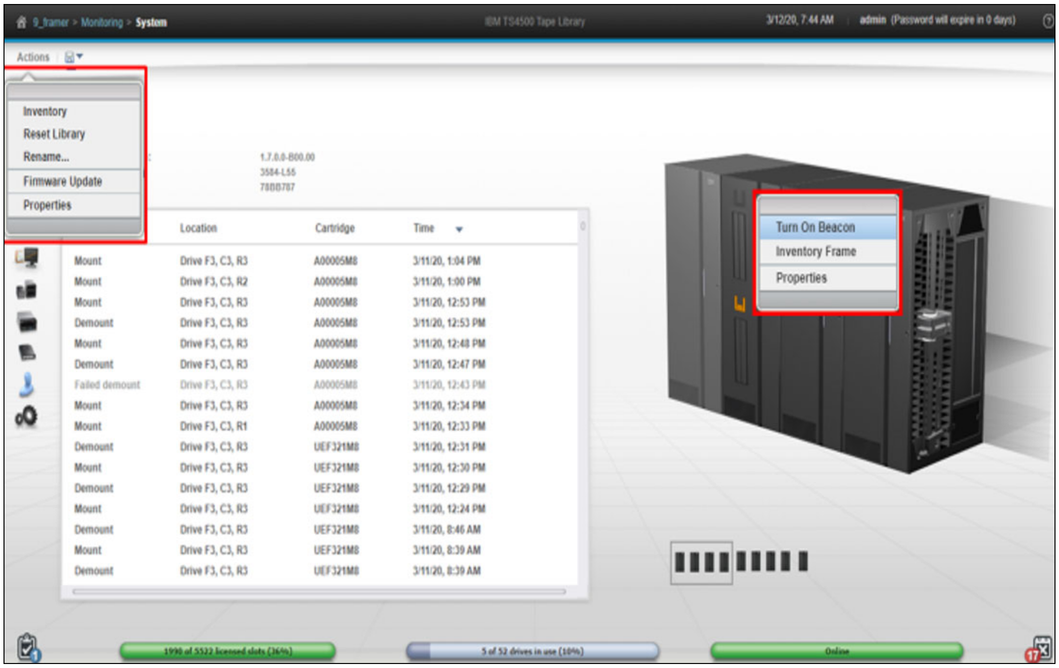

Figure 4-155 Actions menu

You can use certain actions on the System page to perform physical actions on the library remotely. For example, you can open and close I/O station doors from the I/O station Actions menu. You can turn on or turn off the beacon light-emitting diode (LED) from the Actions menu of any frame with an I/O station. You can use the beacon LED to signal remotely to operators the frame that requires attention or the I/O station that contains a specific cartridge. For more information, see “I/O stations” on page 43.

Running an inventory

An inventory operation (see Figure 4-156) includes a check to determine whether each cartridge storage slot in the library is empty or full. It is followed by a scan of the bar code labels. An inventory can be run on a single library frame or an entire library.

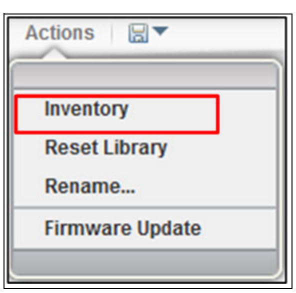

Figure 4-156 Inventory

The TS4500 tape library automatically inventories the library. However, certain conditions necessitate that you manually start an inventory. For example, if an automatic inventory of the media was inconsistent, you might need to run an inventory to reconcile an inconsistency.

Similarly, you might need to start a manual inventory if the host application software indicates that a problem exists with the library inventory.

The TS4500 tape library provides two inventory options, as shown in Figure 4-157.

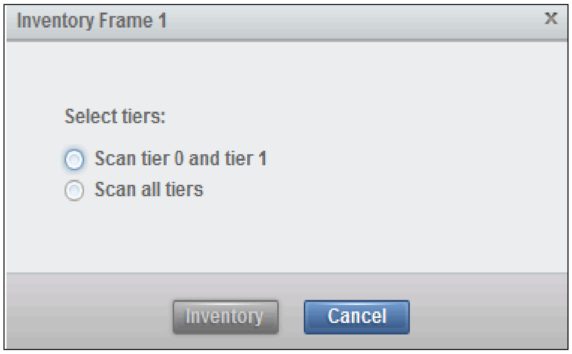

Figure 4-157 Inventory Frame

Select the type of scan that is required:

- Scan tier 0 and tier 1

  Use this option when a faster inventory is preferred. This option scans other tiers only if a discrepancy is identified.

- Scan all tiers

  Use this option when a full library inventory is required. This option is not concurrent, and it can last many hours, depending on the number of cartridges in the library.

> **Note:** The depth of a cartridge location in a HD slot is known as a *tier*. Frame drive side slots are for the cartridges that are immediately accessible in an HD slot (these are Tier 1 cartridges). Behind that is Tier 2, and so on. The maximum tier in an LTO HD slot is Tier 5. The maximum tier in a 3592 HD slot is Tier 4. The single slots on the door-side of a frame are referred to as Tier 0 slots.

When no frame is selected, select **Inventory** from the Actions menu to start an inventory of the library. When a frame is selected, the Inventory option inventories only that frame. Right-click any frame and select **Actions** → **Inventory** to inventory only that frame.

You can inventory a single frame by selecting it and then clicking **Actions** → **Inventory**. If you do not select a frame, all frames are inventoried.

Inventory all tiers

In a high-density (HD) library, a standard inventory is a scan of tier 0 and tier 1. However, at certain times, it is necessary to inventory all tiers. This operation takes more time because it requires moving the cartridges within an HD slot to scan each bar code. For all inventory operations, tier 2 and higher in an HD slot are only scanned when one of the following changes occurs:

- A tier 1 cartridge bar code label changed, which is detected during inventory. Only the slots for which the tier 1 label changed are audited.
- Enough tier 1 bar code labels changed in a column to warrant an audit of the entire column of HD slots.
- An “Inventory all tiers” is requested through the TS4500 management GUI.

Important: We suggest that you limit the manual access of HD slots to the initial bulk loading of the frame because of the length of time that is required to complete a full inventory with an audit.

Typically, the following time frames are required for the TS4500 tape library to audit cartridges:

- Less than 30 seconds for each HD slot
- Up to 10 minutes to inventory a column of full HD LTO slots
- Up to 8 minutes to inventory a column of full HD 3592 slots
- More than 45 minutes to perform a full inventory with an audit on a single HD frame, depending on the frame type

Reset Library

You can reset the TS4500 tape library control system if the library experiences a stopped state. This action (see Figure 4-158) resets the node cards in the library.

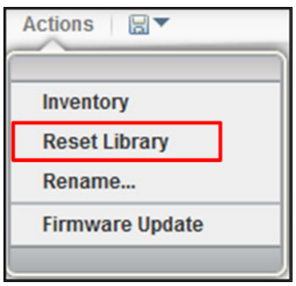

Figure 4-158 Reset Library

A reset can take up to 3 minutes. A reset might require you to reconnect to the web server after the reset completes.

Rename the library

You can rename the library by using the Actions menu and by selecting **Rename**, as shown in Figure 4-159.

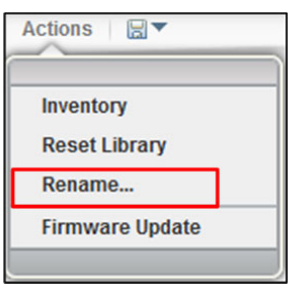

Figure 4-159 Rename the library

Turning on or off the beacon LED

Each frame with an I/O station has a beacon light-emitting diode (LED) that can be turned on or turned off. You can use the beacon LED to identify a frame that is in service or needs service, or to identify a frame for other reasons.

Select a frame with an I/O station. Then, select **Actions** → **Turn On Beacon** or **Actions** → **Turn Off Beacon**, as shown in Figure 4-160.

Figure 4-160 Turning on or turning off the beacon

Updating library firmware

Use the Firmware Update option on the Actions menu to nondisruptively update the library firmware to use the latest library enhancements. You can also use this option to update the drive firmware.

Use the TS4500 tape library to update library firmware without scheduling downtime and without interrupting the job flow or reducing productivity.

A TS4500 tape library firmware upgrade is generally performed by the IBM service support representative (SSR) if the update fixes a field issue or if the update is suggested by IBM. For client-initiated requests, the upgrade is the responsibility of the client. Contact your IBM SSR with any questions about a specific scenario, who is responsible for the upgrade, and whether a charge is associated with the upgrade.

Update procedure

Obtain the library firmware package from IBM Support or from IBM Fix Central at this website and install the package on your workstation.

Complete the following steps to update the firmware:

1. From the System page of the TS4500 management GUI, select **Actions** → **Firmware Update**, as shown in Figure 4-161.

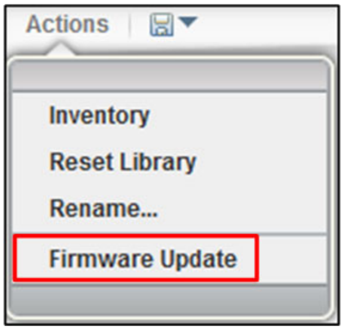

Figure 4-161 Firmware Update

2. On the File Upload window, browse to the library firmware image and click **Open**, as shown in Figure 4-162.

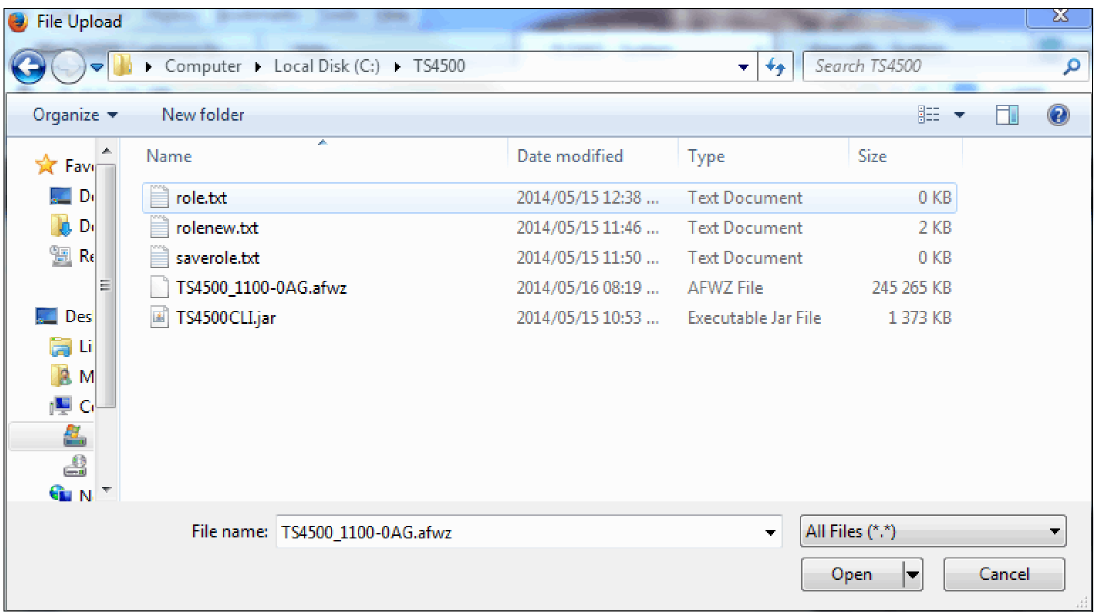

Figure 4-162 File Upload

Note: The library firmware image file is named `TS4500_WXYZ.afwz`, where `W`, `X`, `Y`, and `Z` are different numbers or letters, depending on the firmware version.

3. The Apply Library Firmware window opens. Click **Yes** to continue, as shown in Figure 4-163.

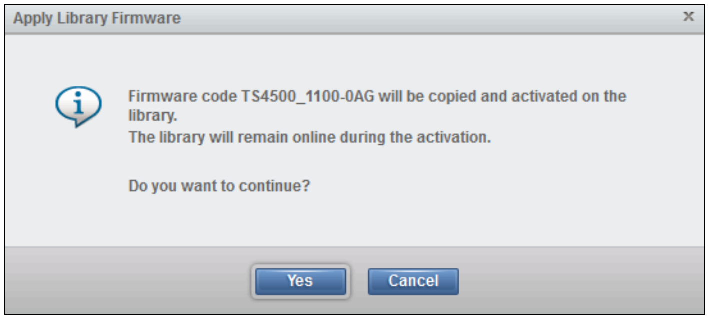

Figure 4-163 Firmware update confirmation

4. The firmware is downloaded to the TS4500 tape library first, and this download can take a few minutes. Then, the update starts and the Tasks pod, which is at the lower-left corner of the System page, indicates that the firmware update task is in progress. Monitor the status of the update from the Tasks page by selecting **Monitoring** → **Tasks**, as shown in Figure 4-164.

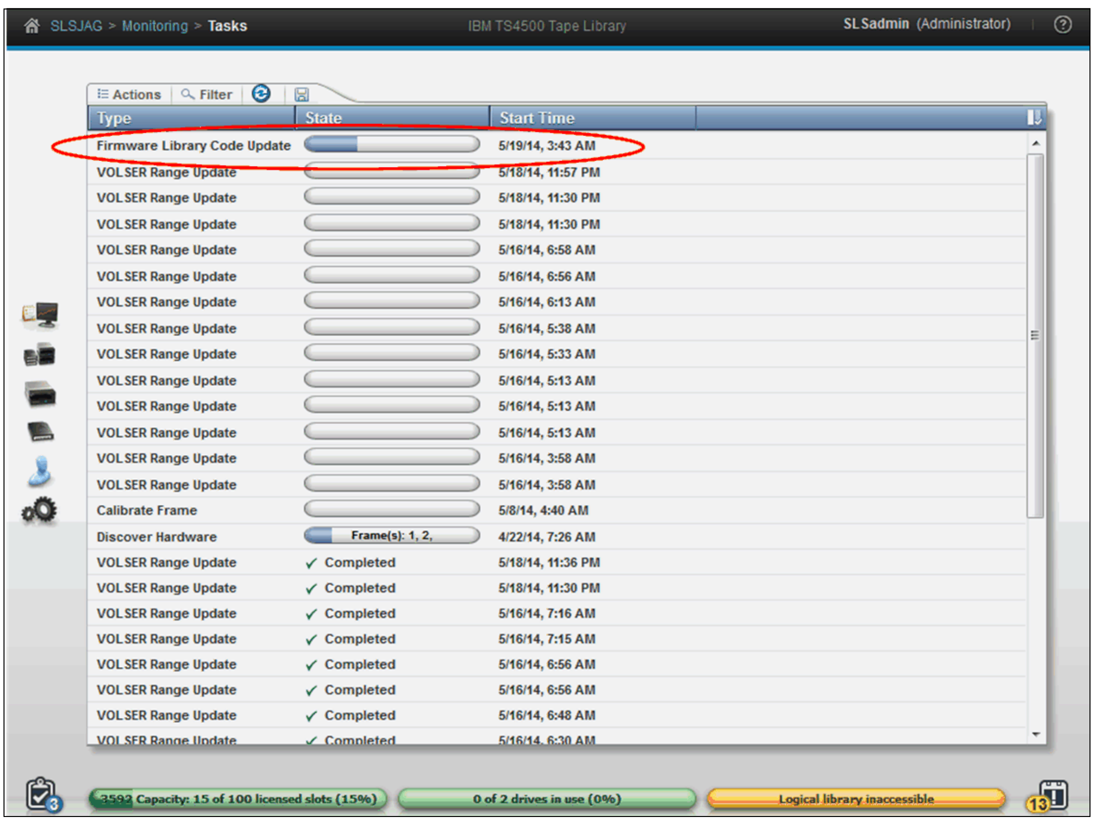

Figure 4-164 Firmware update status

5. You lose connection after the firmware is upgraded. Log back in to the TS4500 tape library and select **Actions** → **Events** to display the status of the update.

Note: If the library firmware update stops, nondisruptively reset the library from the System page by selecting **Actions** → **Reset Library**. Then, try the firmware update again.

Downloading the statistics and logs

Click the **Save** icon (highlighted) on the toolbar to view a list of available downloads, including library logs, component usage statistics, and library configuration information, as shown in Figure 4-165.

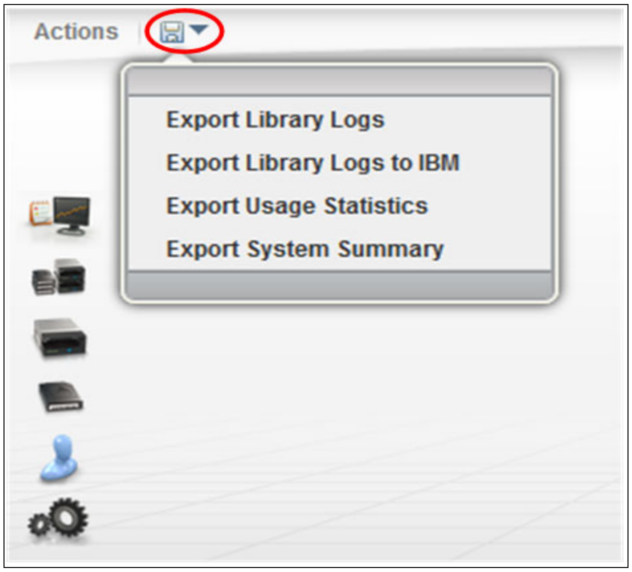

Figure 4-165 Export files

The following files are available for download and export:

- Library logs: This option downloads a `.zip` file that includes event logs, servo logs, nonvolatile random access memory (NVRAM) event logs, and Fatal Exception logs.
- Library Logs to IBM. This option downloads service logs direct to the IBM ECUREP host using the call home functionality on the IMC/TSSC.
- Usage statistics: This option downloads a comma-separated value (`.csv`) file that contains the usage statistics of important library components.
- System summary: This option downloads a `.csv` file that lists the library properties and configuration settings. The downloaded file name uses the following format where *`libraryname`* is the name that you set for the library. The file name format is *`libraryname_YYYYMMDD`*`.csv`.

> **Note:** It is helpful to download this file periodically to compare the library configuration over time or to retain the configuration for auditing purposes.

4.8.2 Events

Events are informational notices, warning alerts, and error alerts that provide more information about the library and library components. Use the Events page to view these events and to monitor their status.

Event categories

Error, warning, and information-only events are displayed on this page. The events are defined in order of severity.

Error

The error icon is a red circle that contains a white X. The error icon is shown in Figure 4-166.

Figure 4-166 Error

Error events are the highest priority events, and they require immediate intervention. They are identified by the error icon (a red circle that contains a white X) in the Events table. Errors indicate a hardware or communication failure that can impair library operations or damage the system.

Warning

The warning icon is a yellow triangle that contains an exclamation mark, as shown in Figure 4-167.

Figure 4-167 Warning

Warning events are the second-highest priority events, and they require attention when time allows. They are identified by the warning icon (a yellow triangle that contains an exclamation mark) in the Events table. Warnings indicate a problem that does not pose an immediate threat, but warnings require resolution to ensure that library operations continue smoothly.

Information

The information icon is a blue balloon that contains the lowercase letter i, as shown in Figure 4-168.

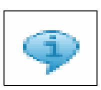

Figure 4-168 Information

Informational events are the lowest-priority events, and they do not require corrective action. They are identified by the information icon (a blue balloon that contains a lowercase letter i) in the Events table. Information-only events provide information about the library or library operations, such as the user login and tape movements.

Monitoring events

The events that are displayed on the Events page are sorted first by severity and second by time. If the Events table is full and a new event is generated, the oldest, inactive events are deleted first. The events table can be limited to show events for 1 day, 1 week, or all. Shown in Figure 4-169 is the Events table.

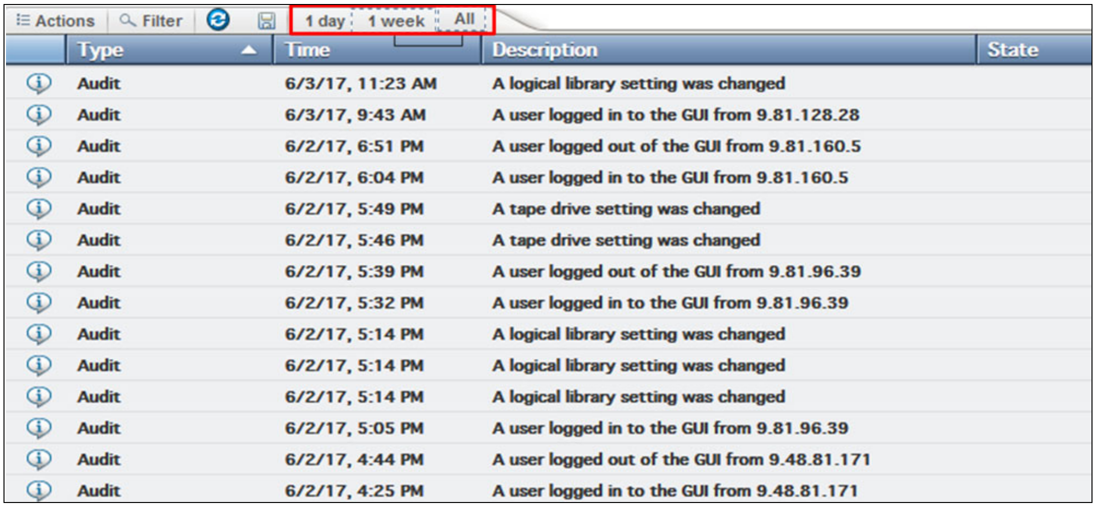

Type	Time	Description	State
Audit	6/3/17, 11:23 AM	A logical library setting was changed	
Audit	6/3/17, 9:43 AM	A user logged in to the GUI from 9.81.128.28	
Audit	6/2/17, 6:51 PM	A user logged out of the GUI from 9.81.160.5	
Audit	6/2/17, 6:04 PM	A user logged in to the GUI from 9.81.160.5	
Audit	6/2/17, 5:49 PM	A tape drive setting was changed	
Audit	6/2/17, 5:46 PM	A tape drive setting was changed	
Audit	6/2/17, 5:39 PM	A user logged out of the GUI from 9.81.96.39	
Audit	6/2/17, 5:32 PM	A user logged in to the GUI from 9.81.96.39	
Audit	6/2/17, 5:14 PM	A logical library setting was changed	
Audit	6/2/17, 5:14 PM	A logical library setting was changed	
Audit	6/2/17, 5:14 PM	A logical library setting was changed	
Audit	6/2/17, 5:05 PM	A user logged in to the GUI from 9.81.96.39	
Audit	6/2/17, 4:44 PM	A user logged out of the GUI from 9.48.81.171	
Audit	6/2/17, 4:25 PM	A user logged in to the GUI from 9.48.81.171	

Figure 4-169 Events

Active events, which are displayed in color in the Events table, are events that are not yet fixed. Inactive events, which are displayed in gray in the Events table, are events that are being addressed or events that are fixed. Use the **Actions** → **Mark Inactive** option to manually mark an event as inactive. Events can also be marked as inactive manually by an IBM SSR.

Only errors and warnings can be marked as inactive. Information-only messages are not alerts, and they cannot be marked as inactive.

Select **Actions** → **Properties** when a specific event is selected to view details about the event. The Event Properties window includes a complete history of the event. If an action can be performed to resolve an error or warning event, a fix procedure is also provided, as shown in Figure 4-170.

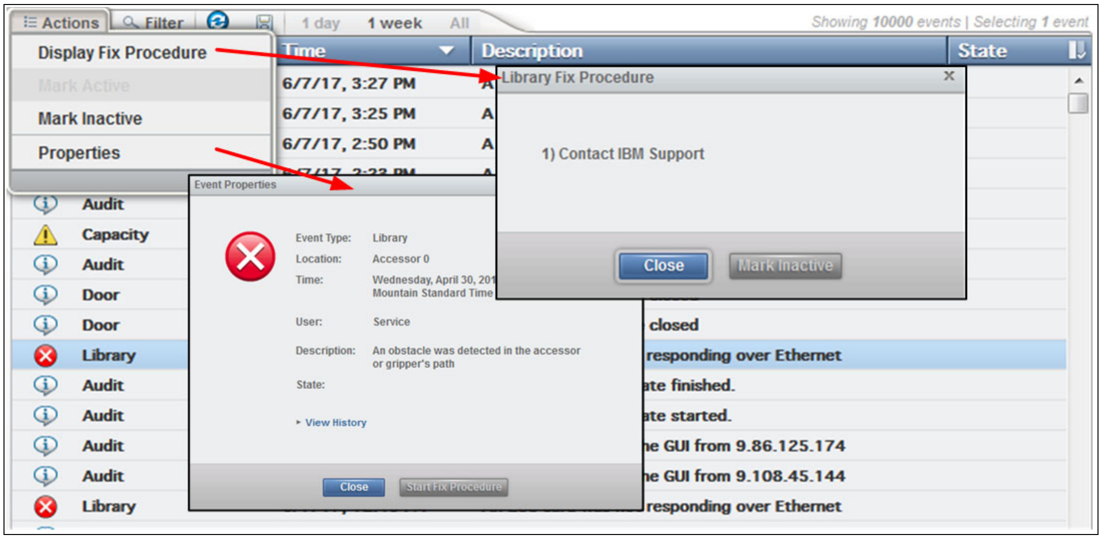

Figure 4-170 Event details

View the status of each event in the State column. For events that require service, the State column shows when a Call Home is generated, when the Call Home is complete, and when the service action is complete.

If SNMP traps are enabled, notification messages are also sent to the SNMP server based on the subscription level (error, warning, or information). Select **Notifications** from the Settings page to manage SNMP traps and subscriptions.

Monitoring tape cartridge moves

Tape cartridge moves are displayed as information events. The default table sort displays all information events last. Sort by location to see a history of moves by drive, I/O station, or storage slot. Sort by time to see the most recent moves, moves on a specific date, or moves at a specific time.

Monitoring when tapes are moved, and where they are moved to and from, can be helpful to view overall library activity. This view can also help you diagnose recurring problems if you use a filtered view of the Events table to display the most recent actions with any component.

Events at a glance

If any events are active, the Events icon displays in the lower-right corner of the System page and the Events page. The number next to the icon indicates the total number of active alerts (both errors and warnings). The color indicates the highest level of the active alerts. If the number is in a red circle, at least one of the indicated alerts is an error (see Figure 4-171). If the number is in a yellow circle, the most severe active alert is a warning.

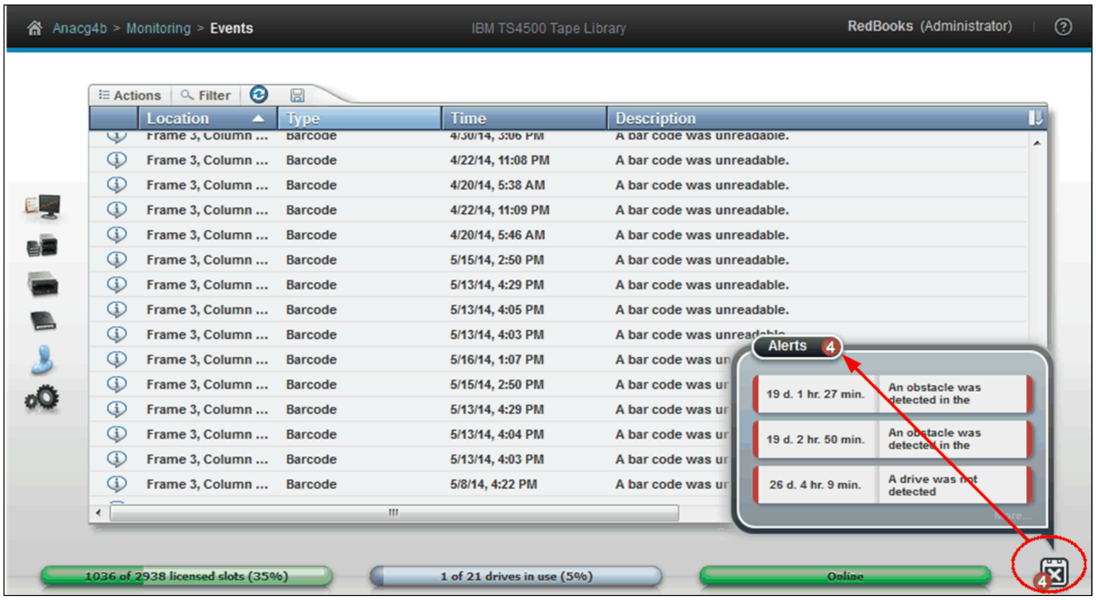

Figure 4-171 Events

Hover over the Events pod to see the location of the event, the relative time since the event occurred in days (d.), hours (hr.), and minutes (min.), and a short description of the event.

Monitoring tasks

The Tasks page displays long-running actions, such as inventory and reset, that run in the background while the library performs other operations. Use this page to monitor the progress of tasks and to view task properties.

The Tasks table displays all active tasks. They are sorted first by state and second by time. Tasks are removed from the table after a certain amount of time when they expire.

Highlight a task in the table, and select **Actions** → **Monitor** to view the status of the task, as shown in Figure 4-172.

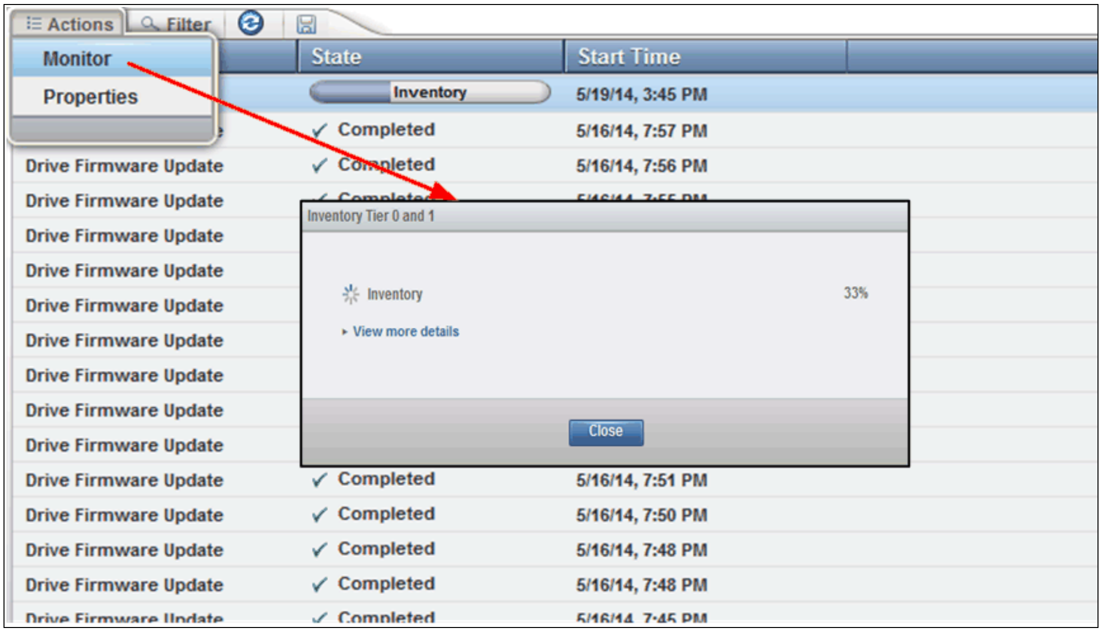

Figure 4-172 Task monitoring

Highlight a task in the table and select **Actions** → **Properties** to view a complete description of the task and the task history, as shown in Figure 4-173.

> **Note:** If a long-running task stops or appears to take more time than necessary, you can reset the library and then start the task again. You can nondisruptively reset the library by selecting **Actions** → **Reset** from the System page to reset the library.

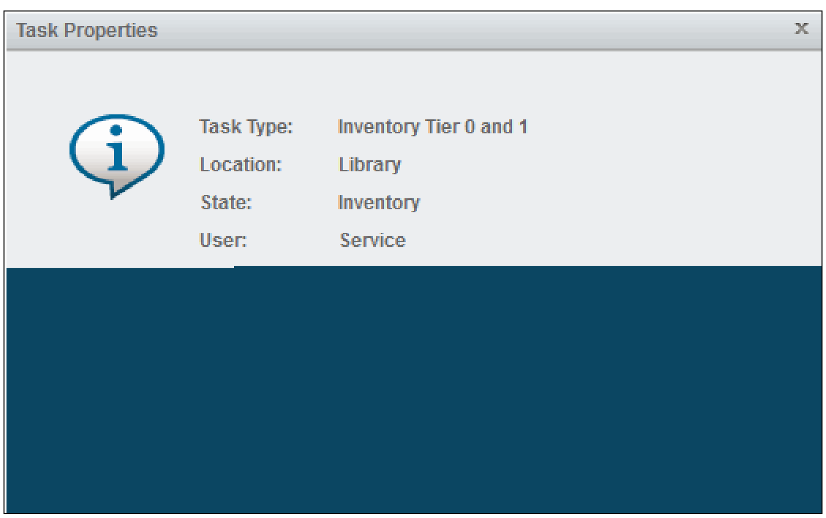

Figure 4-173 Task Properties

The Tasks pod displays in the lower-left corner of the System page when tasks are running. The number next to the Tasks pod indicates the number of tasks. If no icon is displayed, no tasks are active. Hover over the Tasks pod to see the three oldest running tasks that are sorted in the time sequence that they were created, as shown in Figure 4-174.

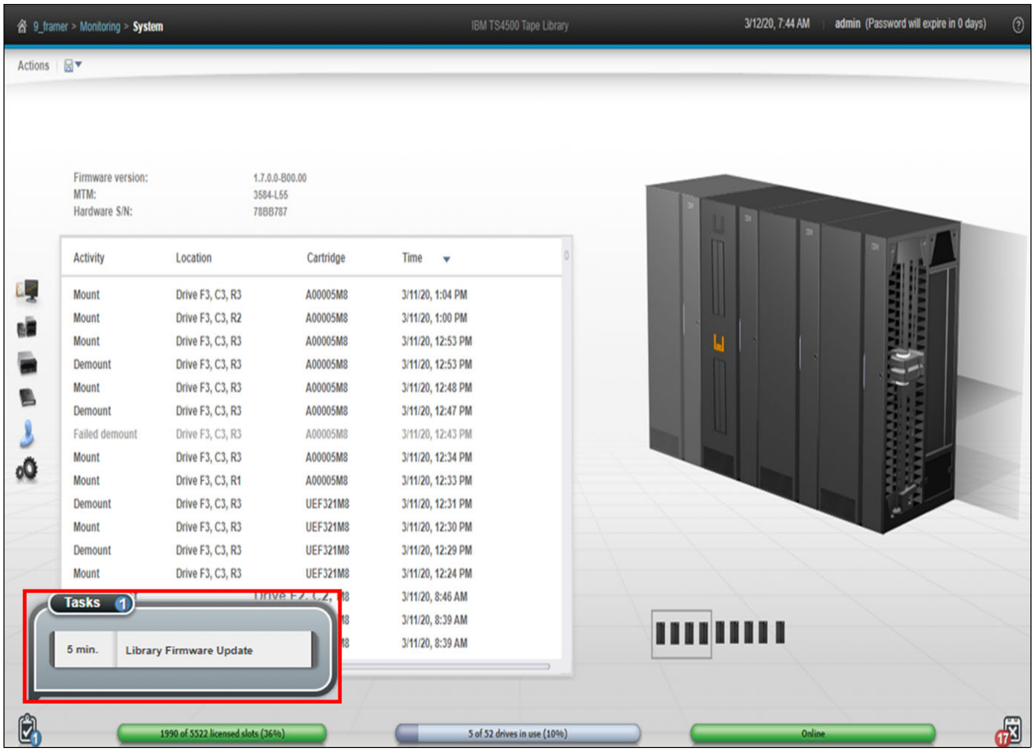

Figure 4-174 Tasks pod

4.9 Tape System Library Manager

Tape System library Manager (TSLM) software provides consolidation and simplification benefits in an IBM TS4500 tape library environment.

TSLM provides a resource management layer between applications, such as IBM Spectrum Protect, and the tape library hardware. TSLM decouples tape resources from applications. This capability simplifies the aggregation and sharing of tape resources.

TSLM can combine the capacity of multiple TS4500 and TS3500 libraries into a single reservoir of tape storage that can be managed from a single point. This capability allows more effective management, monitoring, and reporting of the use of tape storage resources for new and existing TS4500 and TS3500 systems.

TSLM is software that is designed to exceed today’s tape storage management challenges for various clients that are involved in High Performance Computing (HPC) environments, such as oil and gas exploration and genomic analysis, where dozens of tape libraries and hundreds of tape drives must be shared and managed to back up and archive petabytes of data at the lowest cost possible.

The IBM TS4500, IBM TS3500, IBM Tape System Library Manager, and IBM Spectrum Protect are the ideal solutions for these clients.

TSLM provides the following benefits:

- Consolidated, mainframe-class media management services
- Centralized repository, access control, and administration
- Management beyond physical library boundaries:
  - Access to multiple TS3500 or TS4500 tape libraries as a single library image.
  - The libraries can be separate (at SAN distances) or connected in a shuttle complex (TS3500 tape library only).
- Dynamic sharing of resources across heterogeneous application boundaries
- Security features to allow or prevent application access to tapes:
  - Helps to enable a common scratch pool and private pools for every application.
  - Secures the usage and visibility.
- Policy-based drive and cartridge allocation
- Policy-based media lifecycle management
- 3494 emulation

  Emulation of an IBM 3494 library on top of an attached IBM TS3500 or TS4500 tape library.

For more information about TSLM, see *IBM Tape System Library Manager User's Guide*, GA32-2208.

4.10 Remote support

Remote support for the TS4500 tape library involves the use of a Call Home feature to detect and solve problems.

Remote support is available for the TS4500 tape library through its Call Home capability. The *Call Home* feature uses the integrated management console (IMC) to report failures that are detected by the library or a tape drive. Whenever a failure is detected, Call Home sends detailed error information to IBM Support.

The IBM Remote Support Center (RSC) can then prepare an action plan to handle the problem before the SSR travels to the library. The library might also periodically send support information, such as configuration, library, and drive code versions and error logs, to IBM.

The Call Home feature of the TS4500 tape library has three different, but related, capabilities: Problem Call Home, Heartbeat Call Home, and Test Call Home. The TS4500 tape library sends data files that might be helpful to IBM Support Center personnel for all three types of Call Home functions. These data files include library error logs and configuration information, such as the Machine Reported Product Data (MRPD) log. The MRPD file contains the following information about the machine (library):

- Number of frames and drives
- Model and serial number of each frame
- Type and serial number of each drive
- Code version of the library and each drive
- Any machine-detectable features, such as extra I/O stations and capacity expansion

If a Problem Call Home is initiated, the library also sends the tape library logs and drive logs that relate to the problem.

Problem Call Home

The TS4500 tape library or one of its drives detects a problem and the library initiates a Call Home operation. This Call Home operation creates a problem management record (CSP-Ticket) in the IBM Remote Technical Assistance Information Network. This single page of text data can help the IBM Support Center or SSR determine an action plan and a list of parts, which are called field-replaceable units (FRUs).

Heartbeat Call Home

The TS4500 tape library sends the Heartbeat Call Home on a scheduled basis to ensure that the Call Home function operates correctly. By default, the Heartbeat Call Home is sent once a week, 1 hour after a power cycle, and 1 hour after a code update completes. The frequency of Heartbeat Call Home is set on the service menu as shown in Figure 4-175.

A service user is requred to access this option.

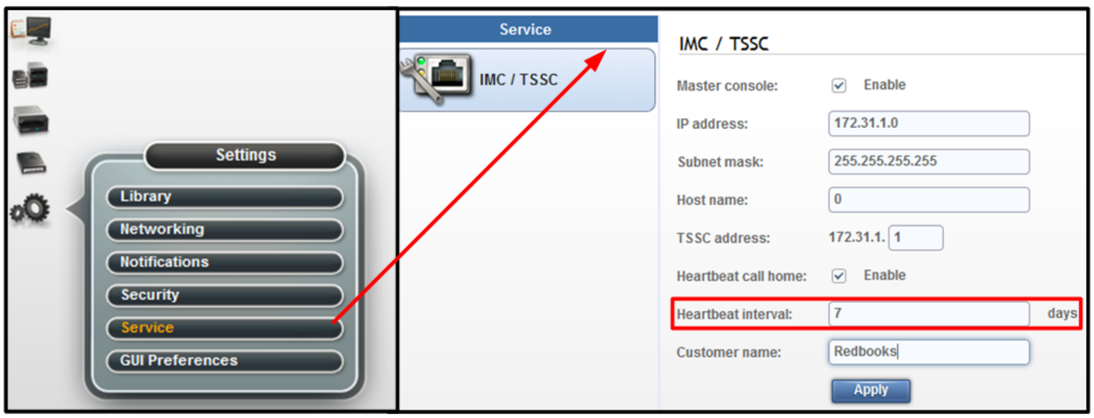

Figure 4-175 Call home settings

Test Call Home

When the SSR services the library, the SSR can issue a Test Call Home operation to RETAIN from the TS4500 management GUI. Through the library, the SSR can include drive dumps in the Test Call Home for analysis. In this way, a drive dump can be accessed by IBM Support Center personnel through the Call Home database.

4.10.1 Remote support through a system console

The integrated management console (IMC) can be used as a service console with the TS4500 tape library. Similar in function to the IBM TotalStorage System Console (TSSC), the IMC comes installed with a set of software tools to help with both the local service and remote support of the attached TS4500 tape library.

Table 4-8 lists the system console features that can be ordered for the library.

Table 4-8 System console and remote support features

Feature code supported models	Supported models	Client setup unit (CSU)	Description
2704	D25, D55, S25 S55, S24, and S54	No	Console expansion 26-port Ethernet switch, rack mount
2715	L25 and L55	No	TSSC Console attachment

The default method that is used by the library for Call Home support is a broadband connection that uses the Electronic Customer Care (ECC) Call Home function through the system console.

To perform an ECC Call Home operation through a system console, the TS4500 tape library sends Call Home information across a private Ethernet connection to the system console. The system console then performs the ECC Call Home operation and sends the information to IBM Remote Technical Assistance Information Network through the Ethernet (broadband) connection.

The IMC is preconfigured for remote support with the TS4500 tape library. The IMC also comes equipped with a RJ45 port that can be used to connect an external network cable for remote support.

Table 4-9 lists the capabilities of remote support with a system console.

Table 4-9 Remote support capabilities

Location	Event support	Support
Client site	Call Home events	► Error initiated ► Heartbeat (regular interval) ► Test
	System console support capability	► Error-initiated problem reporting for up to 43 subsystems ► Staged, error-specific data gathering ► Subsystem and system console heartbeat reporting ► Wellness checking ► Log file storage (daily) ► Code image and documentation repository (from media and RETAIN Fix Distribution Library)
	System console and remote support service tools	► Code image broadcast ► Call Home event log review ► End-of-call completion report
IBM Support	System console remote access	► Authenticated, secure remote access ► Simultaneous call in and Call Home ► Data transmission (Transmission Control Protocol/Internet Protocol (TCP/IP)) supported
	IBM Call Home database	► Always available (24x7) access by IBM Support staff ► Error analysis and search capability

4.10.2 Remote support security

The system console provides Ethernet connectivity through a private internal network. This section describes the security design for remote support of the network-attached devices.

The system console also provides optional Ethernet outbound connectivity through the client's network to the IBM service support system. All inbound communication over this connection is restricted. The system console uses the following protocols to port numbers:

- HTTPS: Port 443
- HTTP: Port 80
- Domain Name System (DNS): Port 53

For outbound and bidirectional data to and from IBM, the system console uses the ports (by default) that are listed in Table 4-10. You can modify the port numbers by using the TS4500 management GUI.

Table 4-10 Default system console ports for outbound and bidirectional data to and from RETAIN

Port	Type of data	Direction	Protocol
67/68	Dynamic Host Configuration Protocol (DHCP)	Outbound	User Datagram Protocol (UDP)
80	HTTP	Bidirectional	Transmission Control Protocol (TCP)
161/162	SNMP	Outbound	TCP
443	Secure Sockets Layer (SSL)	Bidirectional	TCP
1443	IBM Security Key Lifecycle Manager for z/OS (SKLM) server (library-managed encryption (LME) only)	Outbound	TCP
3801	IBM SKLM server (LME only)	Outbound	TCP
389	LDAP	Bidirectional	TCP and UDP
636	LDAPS	Bidirectional	TCP and UDP
443	HTTPS	Bidirectional	TCP

Dial-out security features of the IMC

Dial-out is used by the Call Home feature to send service-related information from the attached systems to the IBM service support system. The following dial-out security properties are available for the attached systems:

- Dial-out is from the client location to the IBM connection point. The IBM service support system does not initiate connections to the attached systems.
- Dial-out through the system console is over an outbound Ethernet connection to the client network. All outbound traffic is limited to HTTP, HTTPS, and Domain Name System (DNS) information. All service-related data is communicated by using HTTPS, and it is encrypted.
- The data that is exchanged between the attached systems and IBM is service-related data. The protocol that is used is specific to this application and not publicly available.
- On the first data exchange of each transmission, IBM validates that the calling system is entitled to service. If the calling system is not validated, it is disconnected.
- The default setting for the Call Home feature is enabled. The Call Home feature can be disabled by an IBM SSR.

None of the client data that is stored on the tape or in memory for the TS4500 tape library is transmitted or accessed in a Call Home session. Call Home is enabled or disabled through the service menu of the TS4500 management GUI.

Dial-in security features of the IMC

Dial-in is used by IBM support to log on to the system console and provide service support. All dial-in connectivity to the system console is through a broadband connection, which restricts all incoming traffic. Separate logon IDs are required for access to each attached system.

Note: The TS4500 tape library does not support dial-in. Only the IMC/TSSC supports dial-in.

The system console supports the following data security requirements when the system console is configured correctly:

- Client data, which is stored on tape or in memory, cannot be transmitted or accessed in remote support sessions.
- Remote dial-in is enabled or disabled by the client by using the TS4500 management GUI. Remote dial-in is disabled, by default. When remote dial-in is enabled, it is enabled for 24 hours, by default.
- Remote dial-in requires a password for access. The password is managed by the client. The default setting for dial-in is to not require a password. A password can be specified by the client and set by the IBM SSR.

Note: Because the IMC and TSSC both run system console code, all of the TSSC functions are integrated into the IMC. The security requirements for the IMC are the same as the security requirements for the TSSC.

Port information for firewall environments

Table 4-11 lists the only ports that are required to be opened on the firewall for environments where the tape configuration is separated from the LAN-attached hosts and web clients by a firewall. All other ports can be closed.

Table 4-11 TS4500 default port information for firewall environments

Function	Port	Direction (from library)	Protocol
Library operations	3494	Bidirectional	TCP
TotalStorage Specialist	80	Inbound	TCP
SNMP traps	161/162	Bidirectional	UDP
Encryption Key Manager	1443	Outbound	SSL
Encryption Key Manager	3801	Outbound	TCP
LDAP	389	Bidirectional	TCP and UDP
LDAPS	636	Bidirectional	TCP and UDP
Kerberos	88	Bidirectional	TCP and UDP
HTTPS and SSL	443	Bidirectional	TCP

Note: The system console uses the following ports:

- HTTPS: Port 443
- HTTP: Port 80
- DNS: Port 53

Port information communications can be initiated by the tape library or by the host. Typically, the library initiates a connection only when it is responding to the host. However, in unsolicited messages, such as statistics notifications and operator interventions, the library initiates a connection through port 3494. If the library needs to make a connection to the host, the library chooses a temporary port and uses that port to make an outbound connection to a 3494 listening port on the host.

When the host has a message to deliver to the library, the host chooses its own ephemeral port by which to make an outbound connection to listening port 3494 on the library manager. The connection is only maintained for the duration that is required to pass a single message, and then the connection is disabled.

Table 4-11 lists the minimum level of connectivity that is required to perform library operations. The following ports can be opened on the firewall, but they are not necessary for full functionality:

- The standard HTTP port, 80, allows inbound communication to the library from the TS4500 management GUI.
- Ports 161 and 162 are the standard ports for sending SNMP traps. SNMP traps can be sent from the TS4500 tape library or from the IMC. The tape library can be configured to send traps to SNMP target machines. In this case, the firewall needs to allow outbound connections from the library from its port 161 to port 162 on the listening SNMP target machine.

Figure 4-176 shows all of the external communication connections to the TS4500 tape library control system.

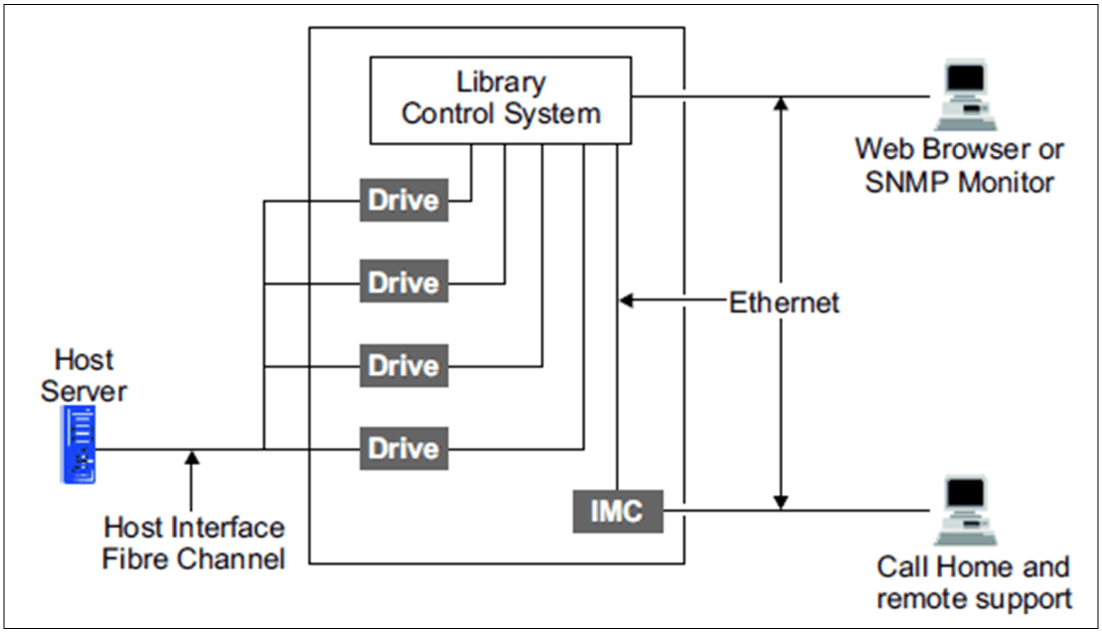

Figure 4-176 External communication connections to the TS4500 tape library control system

The host interface is provided by the drives. The library communicates with the drives through the Library/Drive Interface. The Ethernet port and modem port do not have any direct access to the Library/Drive Interface. All communication between the various interfaces is through the library control system.

All library communication requires explicit support by the library firmware that is running in the library control system. The library firmware does not provide capability for any of the following functions:

- Communicating between the Ethernet port and the TS4500 management GUI
- Communicating between the TS4500 management GUI and the modem port
- Accessing data from tape cartridges through the Library/Drive Interface
- Sending or receiving data from tape cartridges through any port

Security considerations with the TS4500 management GUI

Remote support security for the TS4500 tape library through a system console when you use the management GUI is described.

The management GUI does not allow any access to client data, and it does not allow File Transfer Protocol (FTP) or Telnet type operations. It provides only those functions that are specifically coded in the library firmware. The only files that it can offload are library logs, drive logs, and certain usage and error statistics files. It cannot be used to read or write a client cartridge or otherwise access client data.

The management GUI allows the client to set up an administrator password that is required to use the management GUI to perform any library task. The management GUI also provides several levels of access through various preset roles.

The following list presents potential security concerns when you use the management GUI:

- A management GUI user might move a cartridge from one location to another location within the library, potentially confusing a host application or making a cartridge unavailable by moving it to a different partition.
- A user can reconfigure the library, possibly causing problems at the hosts because of changes in partitioning or device IDs.
- A remote user can update library or drive firmware. However, because the library and drives ignore any firmware that they do not recognize, the only exposure is to loading firmware that is not the latest level.

These security concerns can be addressed by using the password, user role, and authentication features that are provided by the management GUI, and also by enabling SNMP audit logging. When SNMP audit logging is enabled, the library sends notifications when certain events occur in the library.

For more information about security and firewalls that are used to connect the TS4500 and IMC to your local network for Call Home and remote support, see this IBM Support web page.

4.11 IBM Net Promoter Score Feedback

Net Promoter Score (IBM NPS®) is a management tool that can be used to gauge the loyalty of a firm's customer relationships. It serves as an alternative to traditional customer satisfaction research. NPS has been widely adopted with more than two thirds of Fortune 1000 companies using the metric. Every user of the TS4500 GUI is contacted asked after 30 days of use to provide feedback about the TS4500 by way of a “Suggested Task”.

After 30 days, the Suggested Task” is shown (see Figure 4-177 and Figure 4-178).

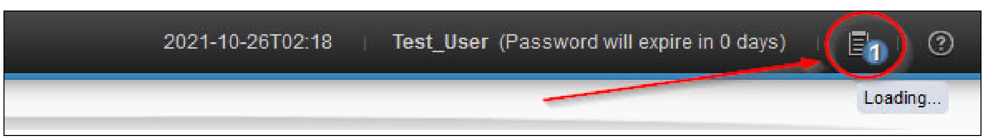

Figure 4-177 New Status icon shows pending task

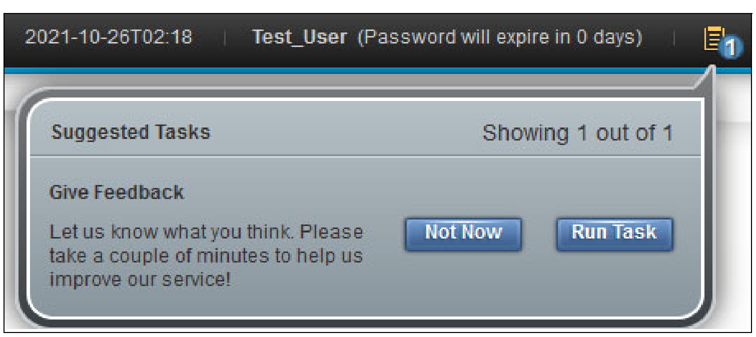

Figure 4-178 Pending feedback task

By clicking **Run Task**, a window opens in which you can start the survey by clicking **Take Survey** or postpone it for 90 days by clicking **Remind me later** (see Figure 4-179).

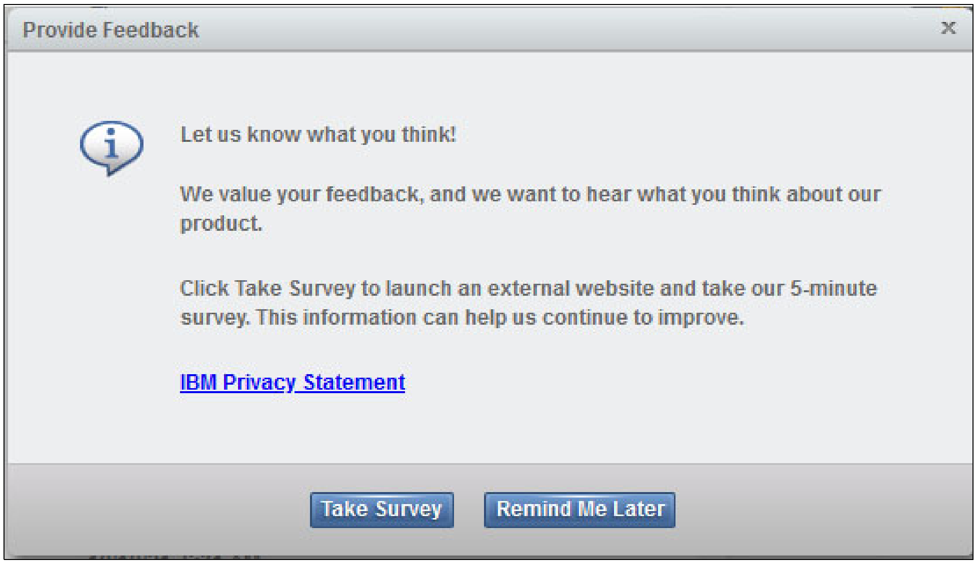

Figure 4-179 Provide feedback dialog

Clicking **Take Survey** opens the Medallia web page where you can provide feedback.

Note: You also can review the IBM Privacy Statement by clicking the corresponding link.

5

Command-line interface

This chapter describes the TS4500 command-line interface (CLI) in the following areas: CLI installation, CLI commands that are categorized by management function, and commands.

This chapter includes the following topics:

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

5.1 TS4500 command-line interface summary

The TS4500 CLI program enables access to TS4500 tape library functions through commands.

The following TS4500 tape library CLI commands are available:

- `assignDataCartridges`
- `assignDriveToLL`
- `batch`
- `bulkAssignDataCartridges`
- `bulkAssignDataCartridgesByLL`
- `cleanDrive`
- `codeUpdate`
- `completeDriveService`
- `continueCloseLibraryVerify`
- `continueLibraryVerify`
- `createBEP`
- `createKeyLabelMapping`
- `createLL`
- `createUser`
- `createVolserRanges`
- `deleteBEP`
- `deleteKeyLabelMapping`
- `deleteLL`
- `deleteUser`
- `deleteVolserRanges`
- `destageDataCartridges`
- `downloadDrivesLog`
- `downloadEvents`
- **`downloadLog`**
- `downloadPropertiesFile`
- `downloadResources`
- `downloadSnapshot`
- `driveCodeUpdate`
- `editKeyLabelMapping`
- `encrypt`
- `getFWVersion`
- `getVIOStatus`**`modifyAdvancedEncSettings`**
- `modifyBEP`
- `modifyFibreChannelSettings`
- `modifyVolserRanges`
- `modifyVolserReporting`
- `moveFromAllDrives`
- `moveFromDrive`
- `moveToDrive`
- `prestageDataCartridges`
- `removeDataCartridges`
- `resetDrive`
- `resetNodeCards`
- `restoreConfiguration`
- `saveConfiguration`
- `setAccessorZones`
- `setAutoEjectCleaningCarts`
- `setDrivePortsId`

- setDriveUse
- setISCSI
- setLibraryTime
- setMacAddress
- setMaximumVIOCartridges
- setNMADetection
- setRolePermissions
- setScannerSpeed
- setSlotOffline
- setSlotOnline
- setSSL
- setUtilThreshold
- showQueuedExports
- startCalibration
- startDriveService
- startDiscoverHW
- startInventory
- startLibraryVerify
- unassignDrive
- version
- viewAccessor
- viewAccessorZones
- viewAdvancedEncryptionSettings
- viewBEP
- viewCleaningCartridges
- viewDataCartridges
- viewDriveDetails
- viewDrivePod
- viewDriveSummary
- viewDriveVPD
- viewFibreChannel
- viewIoStation
- viewISCSI
- viewKeyLabelMapping
- viewLibraryVPD
- viewLogicalLibraries
- viewLogicalLibraryDetails
- viewMacAddress
- viewNodeCards
- viewOfflineComponents
- viewPasswordAndSessionPolicy
- viewRolePermissions
- viewRoles
- viewSnapshots
- viewStorageCapacity
- viewSystemSummary
- viewSystemSummaryDetails
- viewUsers
- viewUtilThreshold
- viewVolserRanges
- viewVolserRangesByLL

CLI commands instruct the system to perform specific data cartridge management, cleaning cartridge management, and service-oriented procedures.

Many of the procedures that are started with CLI commands cannot be started through the TS4500 management GUI.

5.2 Installing the CLI

Download the license and CLI from the IBM TS4500 Command Line Interface web page by clicking **Fix central** → **System Storage** → **Tape systems** → **Tape autoloaders and libraries** → **TS4500 Tape Library (3584)** → **TS4500_CLITool_1.8.0.1**.

You can run the CLI from any directory. However, it is simplest to run it from the directory where the `TS4500CLI.jar` file is stored. By running the CLI from that directory, you can run the CLI commands without entering the path name of the `TS4500CLI.jar` file. However, if you run the CLI from any other directory, you must include the entire path name of the `TS4500CLI.jar` file as a part of the command.

If you run the CLI from the directory where the `TS4500CLI.jar` file is stored, your command looks like the following example:

```
java -jar TS4500CLI.jar -ip [LCC ip] -u [username] -p [password] --[cli_command]
```

If you run the CLI from any other directory, your command looks like the following example:

```
java -jar TS4500CLI.jar "C:\[pathname_of_ts4500cli.jar]\TS4500CLI.jar" -ip [LCC
ip] -u [username] -p [password] --[cli_command]
```

> **Note:** If Secure Sockets Layer (SSL) is enabled on the TS4500 tape library, you must add the **`--ssl`** parameter to the command after the password, as shown in the following example:
>
> ```
> java -jar TS4500CLI.jar -ip [LCC ip] -u [username] -p [password] --ssl
> --[CLI_command]
> ```

For more information about the latest TS4500 CLI commands, see IBM Documentation.

> **Note:** Consider the following points:
>
> - For Library code R8.1, the CLI version 11.8.0.1 or higher is required. The following CLI command is used to check the version:
>
>   ```
>   java -jar TS4500CLI.jar --version
>   CLI GEN 4 Version: 1.8.0.1
>   Build: 07/20/2021
>   ```
>
> - Some commands require service or an admin user to run. If you see error `****ERROR: User name or password incorrect or password expired`, and the user ID and password are correct, the incorrect user ID is being used.
>
>   This list is for V8.1 CLI and might differ from previous versions.

Commands that are categorized by management function

The TS4500 tape library commands are categorized by management function and listed in Table 5-1.

Table 5-1 TS4500 CLI commands that are categorized by management function

Management function	Command usage	Command
Data cartridges	Assigns data cartridges to a logical library.	`assignDataCartridges`
	Assigns drive to a logical library.	`assignDriveToLL`
	Assigns cartridges in bulk instead of one at a time.	`bulkAssignDataCartridges`
	Assigns cartridges from the same logical library in bulk instead of one at a time.	`bulkAssignDataCartridgesByLL`
	Moves cartridges from cartridge cache locations to high-density slots.	`destageDataCartridges`
	Ejects all cartridges.	`moveFromAllDrives`
	Ejects a cartridge from a drive.	`moveFromDrive`
	Mounts a specified cartridge to the specified drive.	`moveToDrive`
	Moves cartridges from high-density slots to cartridge cache locations.	`prestageDataCartridges`
	Removes the data cartridges to and from the I/O station.	`removeDataCartridges`
	Lists high-level information for all of the data cartridges in the library.	`viewDataCartridges`
Cleaning cartridges	Sets the auto eject expired cleaning cartridges flag in the library.	`setAutoEjectCleaningCarts`
	Lists high-level information for all of the cleaning cartridges in the library.	`viewCleaningCartridges`
CLI	Performs actions that are specified within a file.	`batch`
	Creates an encrypted password that can be used with any other CLI command.	`encrypy`
	Used to view the current version of the Web Interface CLI program.	`version`

Management function	Command usage	Command
Fibre Channel Drives	Cleans a specific drive.	`cleanDrive`
	Finish the Drive Service actions. Drive initializes and eventually goes online.	`completeDriveService`
	Sets ports' speed and topology for the specific drive.	`modifyFibreChannelSettings`
	Power cycles a specified drive.	`resetDrive`
	Sets the ports' IDs for a specified drive.	`setDrivePortsId`
	Sets the drive usage configuration.	`setDriveuse`
	Prepare drive for Service and take it offline.	`startDriveService`
	Unassign a specific drive.	`unassignDrive`
	Lists the high-level details of a specific drive.	`viewDriveDetails`
	Show drive statistic similar to the pod in the bottom of the Systems page.	`viewDrivePod`
	Lists the high-level details of all of the drives in the library.	`viewDriveSummary`
	Displays vital product data (VPD) information for all drives.	`viewDriveVPD`
	Displays the Fibre Channel (FC) settings for both ports in each drive.	`viewFibreChannel`
Ethernet Drives	Sets the network configuration settings for a single Ethernet drive or to change the name and alias of the drive.	`setISCSI`
	Sets the MAC address for an Ethernet port or both ports.	`setMacAddress`
	Displays the configuration settings for a specific Ethernet drive.	`viewISCSI`
	Displays the MAC addresses for a drive's Ethernet ports.	`viewMacAddress`
Encryption Key Managers/bar code encryption policy (BEP)	Creates a bar code encryption policy.	`createBEP`
	Edits a bar code encryption policy.	`modifyBEP`
	Deletes a bar code encryption policy.	`deleteBEP`
	Displays a list of all of the VOLSER ranges, showing the BEP settings.	`viewBEP`

Management function	Command usage	Command
Key label mapping	Creates a key label mapping.	`createKeyLabelMapping`
	Edits a key label mapping.	`editKeyLabelMapping`
	Deletes a key label mapping.	`deleteKeyLabelMapping`
	Displays the list of available key labels.	`viewKeyLabelMapping`
Logical libraries	Create a logical library name and associated Drive type (LTO, JAG).	`createLL`
	Delete logical library. No drives must be assigned.	`deleteLL`
	Modifies the advanced encryption settings for a logical library.	`modifyAdvancedEncSettings`
	Modifies the VOLSER reporting value to the host in a six-character format or an eight-character format (`6` or `8`).	`modifyVolserReporting`
	Updates the maximum number of virtual I/O slots for the logical library.	`setMaximumVIOCartridges`
	Sets flag to Show/Hide (`true` or `false`) queued exports.	`showQueuedExports`
	Lists all of the logical libraries.	`viewLogicalLibraries`
	Views the detailed information about a specific logical library.	`viewLogicalLibraryDetails`
	Shows the encryption settings for a certain logical library.	`viewAdvancedEncryptionSettings`
Security	Sets the SSL flag to `Enabled` or `Disabled`.	`setSSL`

Management function	Command usage	Command
Service	Updates the library firmware.	`codeUpdate`
	Used to finish the library verification process for I/O stations.	`continueCloseLibraryVerify`
	Used to perform the second step of the library verification process for I/O stations.	`continueLibraryVerify`
	Downloads a `.csv` file that contains events.	`downloadEvents`
	Downloads the `.zip` file with all of the logs.	`downloadLog`
	Downloads the `.zip` file with all of the logs.	`downloadDrivesLog`
	Downloads the properties file (IBM TotalStorage System Console (TSSC) file).	`downloadPropertiesFile`
	Create a .csv file with information about resources, specified as drive, logical libraries, and so on.	`downloadResources`
	Downloads snapshots.	`downloadSnapshot`
	Updates the drive firmware.	`driveCodeUpdate`
	The command displays the Library firmware version installed.	`getFWVersion`
	Restores the configuration database to the library.	`restoreConfiguration`
	Saves the configuration database to the user's computer.	`saveConfiguration`
	Sets scanner speed.	`setScannerSpeed`
	Sets a specific slot offline.	`setSlotOffline`
	Sets a specific slot online.	`SetSlotOnline`
	Starts the calibration process on selected components of the library.	`startCalibration`
	Start to detect new hardware with the option to delete or keep the current data.	`startDiscoverHW`
	Starts the inventory process in the library.	`startInventory`
	Starts the library verification process for I/O stations.	`startLibraryVerify`
	Displays the list of available snapshots.	`viewSnapshots`

Management function	Command usage	Command
System	Resets all specified node cards.	`resetNodeCards`
	Sets which tape drives and I/O stations are served by which accessor or can also be used to deactivate an accessor.	`setAccessorZones`
	Sets the library date, time, and time zone to that of the server on which the CLI is running.	`setLibraryTime`
	Sets the No Motion Allowed (NMA) detection flag in the library.	`setNMADetection`
	Sets value to Capacity Utilization and Dual Accessor Utilization thresholds.	`setUtilThreshold`
	Shows the accessor's status and usage statistics.	`viewAccessor`
	Displays which tape drives and I/O stations (within sets of frames) are being served by which accessor.	`viewAccessorZones`
	View the actual VIO status	`getVIOStatus`
	Lists the high-level details of all of the cartridges in the I/O station.	`viewIoStation`
	Displays the vital product data (VPD) for the library.	`viewLibraryVPD`
	Displays the information about the node cards.	`viewNodeCards`
	Views offline slots.	`viewOfflineComponents`
	Display Storage slots used for the appropriate media type (Jag / LTO).	`viewStorageCapacity`
	Displays the physical library system summary.	`viewSystemSummary`
	Displays the physical library system summary for the specified frame.	`viewSystemSummaryDetails`
	Show information about Capacity Utilization and Dual Accessor Utilization thresholds.	`viewUtilThreshold`

Management function	Command usage	Command
Users and roles	Create User, Role, and initial password.	`createUser`
	Delete local user. Cannot be the same user who issued the command.	`deleteUser`
	Show information about session definition and password characteristics.	`viewPasswordAndSessionPolicy`
	Revises the existing permissions for a specified role.	`setRolePermissions`
	Displays a list of all of the permissions for a specified role.	`viewRolePermissions`
	Displays a list of all of the defined roles in the library.	`viewRoles`
	Displays a list of all of the users in the library.	`viewUsers`
Volume serial number (VOLSER) ranges	Creates VOLSER ranges.	`createVolserRanges`
	Deletes all of the VOLSER ranges with the file.	`deleteVolserRanges`
	Provides the capability to modify the logical library, media type, and the flag to enable the VOLSER range only for new cartridges.	`modifyVolserRanges`
	Shows all of the VOLSER ranges in the physical library.	`viewVolserRanges`
	Shows all of the VOLSER ranges in a specific logical library.	`viewVolserRangesByLL`

5.3 Commands

The command syntax and examples are described in this section.

If the CLI is run from the directory where the `TS4500CLI.jar` file is stored, the command looks like the following example:

```
java -jar TS4500CLI.jar -ip [LCC ip] -u [username] -p [password] --[cli_command]
```

If the command is run from any other directory, the command looks like the following example:

```
java -jar TS4500CLI.jar "C:\[pathname_of_ts4500cli.jar]\TS4500CLI.jar" -ip [LCC
ip] -u [username] -p [password] --[cli_command]
```

We show only the **[cli_command]** section of the command in the following section. An example of a full command is shown in Example 5-1.

Example 5-1 Command example

```
C:\TS4500>java -jar TS4500CLI.jar -ip 10.1.1.37 -u Redbooks -p RedB00ks
--viewSystemSummary
```

> **Note:** If SSL is enabled in the TS4500 tape library, you must add the **--ssl** parameter to the command after the password. Also, the server must have the correct SSL certificates, as shown in the following example:
>
> `java -jar TS4500CLI.jar -ip [LCC ip] -u [username] -p [password] --ssl --[CLI_command]`

assignDataCartridges

Use the **assignDataCartridges** command to assign or reassign data cartridges to a logical library.

The command uses the following syntax:

`assignDataCartridges filename.txt`

The **assignDataCartridges** command receives a text file with one or more assign operations as input. Each line in the text file is interpreted as an assign operation. In the text file, the data cartridge location can be specified with the format: `[F,C,R,T]`, where `F` = frame, `C` = column, `R` = row, and `T` = tier, as shown in Example 5-2, or the VOLSER, as shown in Example 5-3.

Example 5-2 Using assign_by_location_DataCartridges.txt

```
F1,C2,R8,T0, testltolib1
F2,C2,R8,T0, testltolib2
```

Example 5-3 shows the VOLSER format.

Example 5-3 Using assign_by_volser_DataCartridges.txt

```
A57654L7, testltolib1
A56654L7, testltolib2
```

Example 5-4 shows an example of this command that uses the defined `.txt` files.

Example 5-4 The assignDataCartridges command

```
assignDataCartridges assign_by_location_DataCartridges.txt
assignDataCartridges assign_by_volser_DataCartridges.txt
```

assignDriveToLL

Use the **assignDriveToLL** command to assign drives to a logical library.

The command uses the following syntax:

`assignDriveToLL logicalLibraryname,driveF#C#R#`

Specify the drive locations of the drives, that must be assigned to the logical library. The number sign (#) variables in the syntax have the following values:

F# Frame number of the drive to update
C# Column number of the drive to update
R# Row number of the drive to update

When the **assignDriveToLL** command runs successfully, the system returns output to the display, as shown in Example 5-5.

Example 5-5 The assignDriveToLL command

```
assignDriveToLL TestLib,F1C2R1
The drive was assigned successfully
```

batch

Use the **batch** command to instruct the system to perform multiple commands that are specified in a file.

The command uses the following syntax:

```
batch filename.txt
```

The **batch** command receives a text file with one or more command operations as input. The text file is in the following format:

```
commandName [commandParameters]

next commandName [commandParameters]
```

Each line in the text file is interpreted as a command. When the **batch** command runs successfully, the system returns output that corresponds to the successful execution of each command that is listed in the input file to your display.

Example 5-6 instructs the system to run all of the commands that are contained in the `todays_batch.txt` file.

Example 5-6 The batch command

```
batch todays_batch.txt
```

bulkAssignDataCartridges

Use the **bulkAssignDataCartridges** command to assign cartridges in bulk to a logical library, rather than assigning cartridges one at a time. You can bulk-assign up to 100 cartridges to a common logical library.

The command uses the following syntax:

```
bulkAssignDataCartridges filename.txt, Logical_library_name
```

This command receives a text file with one or more assign operations as input. Each line in the file is interpreted as an assign operation. The input file is in the format `[F,C,R,T] #1` or `VOLSER`, as shown in Example 5-7, where `F` = frame, `C` = column, `R` = row, and `T` = tier.

Example 5-7 Todays_bulk_assigns.txt file

```
F1,C2,R8,T0
F2,C2,R8,T0
F3,C2,R8,T0
```

```
BT1956L7
BT1957L7
```

The logical library name (assignment destination) is the same for all of the entries.

Example 5-8 shows instructing the system to run the assignment of the cartridges that are contained in the `todays_bulk_assigns.txt` file.

Example 5-8 The bulkAssignDataCartridges command

```
bulkAssignDataCartridges todays_bulk_assigns.txt, LibTest
```

bulkAssignDataCartridgesByLL

Use the **`bulkAssignDataCartridgesByLL`** command to assign, in bulk, cartridges from a common logical library source to multiple logical library destinations. The **`bulkAssignDataCartridgesByLL`** command is faster than the **`bulkAssignDataCartridges`** command because the system needs to search in only one logical library source.

The command uses the following syntax:

```
bulkAssignDataCartridgesByLL filename.txt
```

This command receives a text file with one or more assign operations as input. Each line in the file is interpreted as an assign operation. The input file is in the format of `[F,C,R,T] #1`, or `VOLSER` and the destination logical library name, as shown in Example 5-9.

Example 5-9 Input file todays_bulk_assignsLL.txt

```
F1,C2,R8,T0, testltolib1
F2,C2,R8,T0, testltolib2
F3,C2,R8,T0, testltolib1
BT1957L7, testltolib1
BA1955L7, testltolib2
```

Example 5-10 shows an example of this command.

Example 5-10 Command bulkAssignDataCartridgesByLL

```
bulkAssignDataCartridgesByLL todays_bulk_assignsLL.txt
```

cleanDrive

Use the **`cleanDrive`** command to clean a specific drive.

The command uses the following syntax:

```
cleanDrive -f# -c# -r#
```

Each drive location must be delimited by a comma. However, do not insert commas between the `frame #`, `column #`, and `row #`. The number `sign (#)` variables in the syntax have the following values and a space is required between the delimiter:

-f# Frame number of the drive to update
-c# Column number of the drive to update
-r# Row number of the drive to update

When the **cleanDrive** command runs successfully, the system returns output to your display, as shown in Example 5-11.

Example 5-11 The cleanDrive command

```
cleanDrive -f3 -c2 -r1
The drive was cleaned successfully
```

codeUpdate

Use the **codeUpdate** command to update the library firmware. (To update the firmware in a specific drive, use the **driveCodeUpdate** command or use the TS4500 management GUI.)

The command uses the following syntax:

```
codeUpdate filename of firmware image
```

When the **codeUpdate** command runs successfully, the system returns output to your display, as shown in Example 5-12.

Example 5-12 The codeUpdate command

```
codeUpdate TS4500_1801-C00.afwz
Example codeUpdate output
Name: TS4500_1801-C00.afwz
>>>Uploading file...
............................................................
Done
```

completeDriveService

Use the **completeDriveService** command to finish the drive service action. The drive initializes and goes online. The drive location must be specified and set to a service earlier. A task is created and progress can be monitored at the GUI.

The command uses the following syntax:

```
completeDriveService -f# -c# -r#
```

The number `sign (#)` specify the location of the drive, and have the following values:

- **-f#** Frame number of the drive
- **-c#** Column number of the drive
- **-r#** Row number of the drive

When the **completeDriveService** command runs successfully, the system returns output to your display, as shown in Example 5-13.

Example 5-13 The completeDriveService command

```
completeDriveService -f3 -c2 -r1
The drive service has completed
```

continueCloseLibraryVerify

Use the **continueCloseLibraryVerify** command to finish the library verification process for I/O stations. Before you run this command, you must run the **continueLibraryVerify** command and then close all of the I/O station doors. (Library verification is a three-step process in which the following commands must be run in this order: **startLibraryVerify**, **continueLibraryVerify**, and **continueCloseLibraryVerify**.)

The command uses the following syntax:

```
continueCloseLibraryVerify
```

When the **continueCloseLibraryVerify** command runs successfully, the system returns output to your display, as shown in Example 5-14.

Example 5-14 The continueCloseLibraryVerify command

```
continueCloseLibraryVerify
IO doors closed.
Library verify continues now...The test has been completed
successfully
```

continueLibraryVerify

Use the **continueLibraryVerify** command to perform the second step of the library verification process for I/O stations. Before you run this command, you must run the **startLibraryVerify** command and then open all of the I/O station doors. (Library verification is a three-step process in which the following commands must be run in this order: **startLibraryVerify**, **continueLibraryVerify**, and **continueCloseLibraryVerify**.)

The command uses the following syntax:

```
continueLibraryVerify
```

When the **continueLibraryVerify** command runs successfully, the system returns output to your display, as shown in Example 5-15.

Example 5-15 The continueLibraryVerify command

```
continueLibraryVerify
IO doors opened. Library verify continues now.......Done!
Next step: CLOSE all IO doors, and then use "continueCloseLibraryVerify" command
to continue the test.
```

createBEP

Use the **createBEP** command to create the cartridges' bar code encryption policy (BEP). You must run the **viewKeyLabelMapping** command before you run the **createBEP** command to see the key label mapping index number of the cartridges for the encryption policy.

The command uses the following syntax:

```
createBEP VOLSER_start,VOLSER_end,logical_library_name, media_type, klm_index_1,klm_index_2
```

The command includes the following syntax:

VOLSER_start The volume serial number that starts the series of cartridges for the encryption policy.

VOLSER_end The volume serial number that ends the series of cartridges for the encryption policy.

logical_library_name The name of the logical library that contains the series of cartridges for the encryption policy.

media_type The media type of the cartridges for the encryption policy. The following values are valid values for media type: `LTO` or `3592`.

klm_index_1 The key label mapping index number of the cartridges for the encryption policy.

klm_index_2 Optional: This parameter is an optional value to use if you want to create encryption policies for cartridges in a second index.

createBEP When the **createBEP** command runs successfully, the system returns output to your display.

Example 5-16 creates the cartridges' bar code encryption policy for the VOLSER range `TT2000` - `TT3000` in the `test1` logical library for the 3592 cartridges, which are referenced in index 1.

Example 5-16 The createBEP command

```
createBEP TT2000,TT3000,test1,3592,1
The cartridge encryption policy was created successfully
```

createKeyLabelMapping

Use the **createKeyLabelMapping** command to create a key label mapping.

The command uses the following syntax:

```
createKeyLabelMapping keyLabelFrom,keyModeFrom,keyLabelTo
```

Consider the following points regarding the command syntax:

- **keyLabelFrom** is a string that cannot be repeated in the list of key labels. The maximum length of this string is 50 characters.
- **keyModeFrom** is a parameter with the following possible values:
  - **Wrapped-Hash**
  - **Wrapped-Default** (Specifying this value for the **keyModeFrom** parameter disables the **keyLabelTo** parameter, and it is not necessary to specify a string value for the **keyLabelTo** parameter.)
  - **Wrapped-Clear**
- **keyLabelTo** is a string. The maximum length of this string is 50 characters.

When the **createKeyLabelMapping** command runs successfully, the system returns output to your display. In Example 5-17, the system is instructed to create a mapping from the key label `START` to the key label `END` by using the wrapped-hash key mode.

Example 5-17 The createKeyLabelMapping command

```
createKeyLabelMapping START,Wrapped-Hash,END
The Key Label Mapping was created successfully
```

createLL

Use the **createLL** command to create a logical library. Library name and the Drive Media Type associated with the Library to be specified.

The command includes the following parameters:

- **Name**: Name of the new logical library
- **Media** type: LTO or JAG

The command uses the following syntax:

```
createLL NewLibname,DriveMediaType
```

When the **createLL** command runs successfully, the system returns output to the display, as shown in Example 5-18

Example 5-18 The createLL command

```
createLL NewLib,JAG
The logical library was created successfully
```

createUser

Use the **createUser** command to create user names, roles, and a temporary password, if required.

The command uses the following parameters:

- **name**: User name
- **role**: Access level, Admin, Service, Superuser, Monitor
- **sendToEmail**: Email notification yes/no
- **tempPass**: Set temporary password yes/no

Consider the following points:

- The email parameter is required if **–sendToEmail** is 'yes' and optional if **–sendToEmail** is 'no'.
- The **–tempPass** parameter is required if **–sendToEmail** is 'no' and not allowed if **–sendToEmail** is 'yes'.
- The **–name** and **–role** parameters always are required.
- Only the Administrator and Super User are authorized to use this command.

The command uses the following syntax:

```
createUser -name <name> -role <role>  -sendToEmail [yes {-email <email> | no
-tempPass <password> }
```

When the **createUser** command runs successfully, the system returns output to the display, as shown in Example 5-19.

Example 5-19 The createUser command

```
createUser –name RobRoy123 –role Administrator –sendToEmail no –tempPass
L0gM3InN0w
createUser –name RobRoy123 –role Service –sendToEmail yes –email
RobRoy123@mycompany.com
User RobRoy123 was created successfully.
```

createVolserRanges

Use the **createVolserRanges** command to create one or more VOLSER ranges. This task is a long-running task. The CLI shows the percentage of task completion as the task progresses.

The command uses the following syntax:

```
createVolserRanges filename.txt
```

The **createVolserRanges** command receives a text file with one or more VOLSER ranges as input. Each line of the input file is interpreted as a VOLSER range. The input file is in the following format:

```
startRange,endRange,LLName,mediaType,Enable
```

The following parameter information is required:

- **startRange** is the starting VOLSER in the range to create.
- **endRange** is the ending VOLSER in the range to create.
- **LLName** is the name of the logical library.
- **mediaType** is `LTO` or `3592`.
- **Enable**:
  - **TRUE** = Enable this range for new cartridges only.
  - **FALSE** = Do not enable this range for new cartridges only.

While each line of the input file runs successfully and each VOLSER range is created, the system returns the results to your display.

Example 5-20 is an example of the **createVolserRanges** input file.

Example 5-20 The volser_ranges.txt input file

```
UAA9RH,UBB9RH,Library1,3592,FALSE
```

Example 5-21 uses the input file to instruct the system to create the VOLSER range that starts with `UAA9RH` and ends with `UBB9RH` in logical library `1` for the 3592 frame. This VOLSER range is not enabled for new cartridges.

Example 5-21 The createVolserRanges command

```
createVolserRanges volser_ranges.txt
The Volser Range was created successfully
Done
```

deleteBEP

Use the **deleteBEP** command to delete a cartridge bar code encryption policy. You must run the **viewBEP** command before you run this command to see the list of valid indexes.

The command uses the following syntax:

```
deleteBEP index
```

Example 5-22 shows a command that instructs the system to delete the BEP for index 1. It is necessary to run the **viewBEP** command before you run the **deleteBEP** command to see the index number.

Example 5-22 The deleteBEP command

```
deleteBEP 1
The cartridge encryption policy was deleted successfully
```

deleteKeyLabelMapping

Use the **deleteKeyLabelMapping** command to delete a key label mapping. You must run the **viewKeyLabelMapping** command before you run this command to see the index of the key label mapping that you want to delete.

The command uses the following syntax:

```
deleteKeyLabelMapping index
```

When the **deleteKeyLabelMapping** command runs successfully, the system returns output to your display, as shown in Example 5-23.

Example 5-23 The deleteKeyLabelMapping command

```
deleteKeyLabelMapping 1
The Key Label Mapping was deleted successfully
```

deleteLL

The command **deleteLL** deletes a specific logical library

The command uses the following syntax:

```
deleteLL Libraryname
```

The **Name** parameter name of the logical library

When the **deleteLL** command runs successfully, the system returns output to the display, as shown in Example 5-24.

Example 5-24 The deleteLL command

```
deleteLL TestLib
The logical library was deleted successfully
```

deleteUser

The **deleteUser** command delete local user accounts. After deletion, the account cannot be used for authentication to the TS4500 management GUI or CLI.

The command uses the following syntax:

```
deleteUser -username
```

Consider the following points:

- This action behaves like the matching GUI action and follows all its rules. This behavior includes not being able to delete the user account that initiated this command.
- Only the Administrator and Super User are authorized to use this command.

When the **deleteUser** command runs successfully, the system returns output to the display, as shown in Example 5-25.

Example 5-25 The deleteUser command

```
deleteUser test_admin
The user was deleted successfully
```

deleteVolserRanges

Use the **deleteVolserRanges** command to delete VOLSER ranges. This task is a long-running task.

The command uses the following syntax:

```
deleteVolserRanges filename.txt
```

The **deleteVolserRanges** command receives a text file with one or more VOLSER ranges as input. Each line of the input file is interpreted as a VOLSER range to delete. The input file is in the following format: `LLName,startRange,endRange`. Consider the following points regarding the input file:

- **LLName** is the name of the logical library.
- **startRange** is the starting VOLSER in the range to delete.
- **endRange** is the ending VOLSER in the range to delete.

Example 5-26 shows an example of a **deleteVolserRanges** input file. The example input file instructs the system to delete the VOLSER that starts with `TUU9RH` and ends with `TXX9RH` from logical library `1`.

Example 5-26 The ranges_to_delete.txt file

```
Library 1,TUU9RH,TXX9RH
```

Example 5-27 show the command. As each line of the input file runs successfully and each VOLSER range is deleted, the system returns the results to your display.

Example 5-27 The deleteVolserRanges command

```
deleteVolserRanges ranges_to_delete.txt
The Volser Range was removed successfully
Done
```

destageDataCartridges

Use the **destageDataCartridges** command to move cartridges from their cartridge cache locations as specified in a text file or by using a list of VOLSERs to a high-density slot. This command requires that all data cartridges that are being destaged are in tier 0.

The command uses the following syntax:

```
destageDataCartridges filename.txt or
destageDataCartridges [VOLSER1],[VOLSER2],[VOLSERX],...
```

The **destageDataCartridges** function can use a text file with one or more destage operations as input. The `filename.txt` variable specifies the file name of the input file. The text file is in one of the following formats:

- From a storage location: `[F,C,R,T]`
- VOLSER: `[VOLSER]`

Each line in the text file is interpreted as a prestage operation. In the text file, the data cartridge location can be specified with the `[F,C,R,T]` location or the `VOLSER`.

> **Note:** All source cartridges that are being destaged must be in tier 0 (T0).

Example 5-28 shows a **destageDataCartridges** input file with the F,C,R,T reference.

Example 5-28 Input file with F,C,R,T reference

```
F1,C2,R8,T0
F1,C3,R18,T5
```

Example 5-29 shows a **destageDataCartridges** input file with a VOLSER reference.

Example 5-29 Input file with VOLSER reference

```
VOLSER UAA9RHL7
VOLSER UBB9RHL7
```

Example 5-30 shows an example of the **destageDataCartridges** command. When the **destageDataCartridges** command runs successfully, the system returns output to your display. The **destageDataCartridges** command instructs the system to destage the data cartridges from the correct library location based on the information in the `destage.txt` input file.

Example 5-30 The destageDataCartridges command

```
destageDataCartridges destage.txt
The cartridge [location] was moved successfully. Next element...
...
Done
```

downloadEvents

Use the **downloadEvents** command to export detailed information about all error, warning, and informational events saved in the library. It is sorted first by severity and second by time. This information is saved as a `.csv` file in the same directory in which the command was run.

The command uses the following syntax:

```
downloadEvents
```

When the **downloadEvents** command runs successfully, the system returns output to the display, as shown in Example 5-31.

Example 5-31 The downloadEvents command

```
downloadEvents
Downloading...536220/536220 bytes
The events file: TS4500_Events_2018-10-02_12.14.44.csv has been downloaded.
```

downloadLog

Use the **downloadLog** command to download a .zip file with the logs from all node cards. This command is useful for troubleshooting.

The command uses the following syntax:

```
downloadLog
```

The **downloadLog** command does not have any required parameters. It downloads all of the logs for all of the libraries. When the command runs successfully, the system returns output to your display, as shown in Example 5-32.

Example 5-32 The downloadLog command

```
downloadLog
Downloading... [completed] / [totalSize] has been downloaded.
The log file: "TS4500_FWLOGS_[yyyy-MM-dd_hh.mm.ss].zip has been downloaded.
```

downloadDrivesLog

Use the **downloadDrivesLog** command to download a .zip file containing the logs from the specified (or all) drives. This command is useful for troubleshooting.

The command uses the following syntax:

```
downloadDrivesLog
driveCodeUpdate F#C#R#,F#C#R#, / ALL
```

Specify the drive locations of the drives, that you want to retrieve logs from. You can select to download all drives at the same time. Each drive location must be delimited by a comma. However, do not insert commas between the `frame #`, `column #`, and `row #`. The number sign (#) variables in the syntax have the following values:

F# Frame number of the drive to update
C# Column number of the drive to update
R# Row number of the drive to update
All All drives

The **downloadDrivesLog** command does not have any required parameters. It downloads the logs for all the drives requested, or using the **All** option downloads from all drives. When the command runs successfully, the system returns output to your display, as shown in Example 5-33.

Example 5-33 The downloadDrivesLog command

```
downloadDrivesLog F1C1R1,F1C2R2
Gathering logs...
Downloading...3638135/3638135 bytes
The drivelog file: TS4500_DRIVELOGS_2017-06-06_11.41.58.zip has been downloaded
```

downloadPropertiesFile

Use the **downloadPropertiesFile** command to download the properties file (IBM TotalStorage System Console (TSSC) file).

The command uses the following syntax:

```
downloadPropertiesFile
```

The **downloadPropertiesFile** does not have any required parameters. When the command runs successfully, the system returns output to your display, as shown in Example 5-34.

Example 5-34 The downloadPropertiesFile command

```
downloadPropertiesFile
Downloading... [completed] / [totalSize] has been downloaded.
The file: ""LIBLG_01_VP_[yyyy-MM-dd_hh.mm.ss] has been downloaded.
```

downloadResources

Use the **downloadResoruces** command to export detailed information about all resources of a specific type for this library. This information is saved as a `.csv` file. The format of this file matches the exported table information from the management GUI.

The command uses the following syntax:

```
downloadResources
<LogicalLibraries,Drives,FibreChannelPorts,iSCSIPorts,Slots,Cartridges,VOLSERRange
s,Users,Roles,ManagementEthernetPorts>
```

The file that is generated by microcode is saved to the working directory where the CLI is running. The name of this file should use the following format:

```
<lib name>_<lib S/N>_<resource_type>_<timestamp>.csv
```

In the case where “All” is selected, the filename is `<library_name>_LibrarySummary_<timestamp>.csv`. Otherwise, `<resource_type>` matches what is put into the command, as shown in Example 5-35.

Example 5-35 The downloadResources command

```
downloadResources LogicalLibraries
Downloading...372/372 bytes
The resources file: PFETS4500MZ_78AA004_LogicalLibraries_20181002T033045.csv has
been downloaded.
```

Figure 5-1 shows output from the `.csv` file.

Name	Media type	Cartridges	Virtual I/O Cartridges	Encryption	Drives	Maximum Cartridges
"FVT_NEW"	"3592"	"4"	"0"	"None"	"1"	"3938"
"55F_temp1"	"3592"	"3"	"0"	"None"	"1"	"50"
"LTFSEE_Saitama"	"3592"	"18"	"0"	"None"	"5"	"3938"
"LTFSEE_Mikasa"	"3592"	"27"	"6"	"Library managed (Barcode)"	"5"	"3938"
"FVT"	"3592"	"9"	"6"	"None"	"2"	"10"
"55E_Tmp"	"3592"	"16"	"2"	"None"	"4"	"3938"
"rao"	"3592"	"15"	"6"	"Library managed (Internal selective)"	"0"	"1000"
"Jag4_test"	"3592"	"2"	"0"	"System managed"	"2"	"3938"
"JAG6_FC2"	"3592"	"0"	"0"	"None"	"2"	"3938"
"TEST"	"3592"	"0"	"0"	"System managed"	"2"	"3938"
"jag6e_test"	"3592"	"3"	"0"	"Library managed (Barcode)"	"1"	"3938"
"LTFSEE_Shibuya"	"3592"	"4"	"0"	"None"	"2"	"3938"
"LTFSEE_Tora"	"3592"	"9"	"0"	"Library managed (Barcode)"	"2"	"3938"

Figure 5-1 downloadResources .csv file example

downloadSnapshot

Use the **downloadSnapshot** command to display a list of all available snapshots in the library. You must run the **viewSnapshots** command before you run the **downloadSnapshot** command to display a current list of indexes.

The command uses the following syntax:

```
downloadSnapshot index
```

The **index** is the index number of the snapshot to download. (The index numbers are displayed with the **viewSnapshots** command.)

The snapshot file is downloaded to the same directory where the CLI is running. When the **downloadSnapshot** command runs successfully, the system returns output to your display, as shown in Example 5-36.

Example 5-36 The downloadSnapshot command

```
downloadSnapshot 1
Downloading... [completed] / [totalSize] has been downloaded.
..
The snapshot file: "TS4500_Snapshot_01122013.zip has been downloaded.
```

driveCodeUpdate

Use the **driveCodeUpdate** command to update the firmware in a specified drive or specified list of drives.

The command uses the following syntax:

```
driveCodeUpdate Filename of firmware image -l F#C#R#, F#C#R# --reset reset_option
```

The list of drives that you want to update must be preceded by an **-l**. Specify the drive locations of the drives that you want to update. Each drive location must be delimited by a comma. However, do not insert commas between the `frame #`, `column #`, and `row #`. The `number sign (#)` variables in the syntax have the following values:

F# Frame number of the drive to update
C# Column number of the drive to update
R# Row number of the drive to update
--reset Reset method for the drive after the code is updated

The following values are the valid values for the **reset** parameter:

IMMEDIATE The reset occurs immediately after the code is on the drive. The drive must be empty before the update. If the drive has a tape before you perform an `IMMEDIATE` reset, run the **moveFromDrive** command to eject a cartridge from the drive.

UNLOAD The reset occurs after the drive is unloaded.

MANUAL A manual reset must be performed to activate the code.

When the command runs successfully, the **driveCodeUpdate** command instructs the system to update the drives in the list. The system returns output to your display. Example 5-37 on page 349 shows that the system is instructed to update the drives in frame `1`, column `4`, row `2`, and frame `1`, column `4`, row `3` by using the code update that is contained in the `*.fmrz` file.

Example 5-37 The driveCodeUpdate command

```
driveCodeUpdate LTO9_N9BO.ssp_fh.fmrz -l F1C4R2,F1C4R3 --reset IMMEDIATE
Name: LTO9_N9BO.ssp_fh.fmrz

These drives will be updated: F1C4R2,F1C4R3
>>>Uploading file..
.....
Done.
The code update has started
----------------------------------------
```

editKeyLabelMapping

Use the **editKeyLabelMapping** command to edit a key label mapping. You must run the **viewKeyLabelMapping** command before you run this command to see the index of the key label mapping that you want to edit. The command uses the following syntax:

```
editKeyLabelMapping index,keyLabelFrom,keyModeFrom,keyLabelTo
```

Consider the following points regarding the command syntax:

- **index** is the index number of the key label to edit. You must run the **viewKeyLabelMapping** command to see the index number.
- **keyLabelFrom** is a string that cannot repeat in the list of key labels. The maximum length of this string is 50 characters.
- **keyModeFrom** is a parameter with the following possible values:
  - **Wrapped-Hash**
  - **Wrapped-Default**: Specifying this value for the keyModeFrom parameter disables the keyLabelTo parameter. Also, you do not need to specify a string value for the keyLabelTo parameter.
  - **Wrapped-Clear**
  - **Direct-Default-Set**
  - **Direct-Specific**
- **keyLabelTo** is a string. The maximum length of this string is 50 characters.

When the **editKeyLabelMapping** command runs successfully, the system returns output to your display, as shown in Example 5-38. The **editKeyLabelMapping** command instructs the system to edit the key label mapping for index `1`. The revised mapping is from the key label `START` to the key label `END` by using the `Wrapped-Hash` key mode.

Example 5-38 The editKeyLabelMapping command

```
editKeyLabelMapping 1,START,Wrapped-Hash,END
The Key Label Mapping was edited successfully
```

encrypt

Use the **encrypt** command to create an encrypted password that can be used with any other CLI command. When an encrypted password is used in a command, use **-ep** instead of **-p** before the password.

The command uses the following syntax:

```
encrypt [PlainTextPassword] -ssl
```

Consider the following points:

- **<PlainTextPassword>** is a required field and is the plain text password that should be encrypted.
- **-ssl** is required if SSL (HTTPS) is enabled.

The **encrypt** command is shown in Example 5-39.

Example 5-39 The encrypt command

```
encrypt AdminPW1
xEeUtEr9
```

In this example, the command encrypted the password `AdminPW1` and the TS4500 responded with an encrypted version of this password; therefore, any command after this can use the encrypted version. As shown in Example 5-40, the encrypted password `xEeUtEr9` is used with the **-ep** parameter to run the **viewSystemSummary** command.

Example 5-40 Using the encrypted PW

```
java -jar TS4500CLI.jar -ip 10.129.251.127 -u admin -ep xEeUtEr9
--viewSystemSummary
```

getFWVersion

Use the **getFWVersion** command to return the FW version that is installed in the library. The command uses the following syntax:

```
getFWVersion
```

When the **getFWVersion** command runs successfully, the system returns output to the display that is shown in Example 5-41.

Example 5-41 Using the getFWVersion

```
getFWVersion
Firmware Version: 1.8.0.1-C00
```

getVIOStatus

Use the getVIOStatus to view the VIO status (Enabled or Disabled). The command uses the following syntax:

```
getVIOStatus
```

When the **getVIOStatus** command runs successfully, the system returns output to the display that is shown in Example 5-42.

Example 5-42 Using the getVIOStatus

```
VIO flag is : Enabled
```

modifyAdvancedEncSettings

Use the **modifyAdvancedEncSettings** command to make the advanced encryption settings more or less restrictive to data stored in a library.

The command uses the following syntax:

```
modifyAdvancedEncSettings <Logical Library Name>, <Advanced Method> [TRUE /
FALSE], <Advanced Policy>, <Density code>, <keypath>
```

The following values are valid for **Advanced Policy**:

- No advanced setting = 0
- Do not encrypt if no policy = 1
- Encrypt if no policy = 2
- Policy required = 3
- Never encrypt (policy override) = 4
- Always encrypt (policy override) = 5
- Internal label: selective encryption = 6
- Internal label: encrypt all = 7

The following values are valid for the **Density code** setting:

- No advanced setting = 0
- Show encryption = 1
- Mask encryption = 2

The following values are valid for the **keypath** setting:

- No advanced setting = 1
- System = 2
- Application (IBM) = 3
- Application (T10) = 4
- Library = 6

As shown in Example 5-43, the **modifyAdvancedEncSettings** operation modifies the encryption settings for `Library1` by using the advanced policy that the system must encrypt the data if the data is not attached to an encryption policy. The density code of `2` instructs the system to mask the encryption. The key path of `2` is an instruction that the path to the encryption key to encrypt the incoming data is stored in the system, rather than in the specific library or specific application.

Example 5-43 The modifyAdvancedEncSettings command

```
modifyAdvancedEncSettings Library1,TRUE,2,2,2
The advanced encryption settings for Logical Library Library1 were updated
successfully
```

modifyBEP

Use the **modifyBEP** command to edit the cartridges' bar code encryption policy. You must run the **viewKeyLabelMapping** and the **viewBEP** commands before you run this command to see the list of valid indexes.

The command uses the following syntax:

```
modifyBEP index -KLM1 # -BEP true/false -KLM2 #
```

Consider the following points regarding the command's syntax:

index The index number of the cartridges' encryption policy to edit. You must run the **viewBEP** command to see a current list of valid index numbers.

-KLM1 The index number of the key label mapping that contains the encryption policy to edit. You must run the **viewKeyLabelMapping** command to see a current list of valid index numbers.

-KLM2 Optional: An optional parameter to use if you want to edit encryption policies in a second key label mapping index.

-BEP A flag to enable BEP (value is `true` or `false`).

When the **modifyBEP** command runs successfully, the system returns output to your display. Example 5-44 shows the command that updates the cartridges' bar code encryption policy in BEP index 1 and key label mapping index 3 with BEP not enabled.

Example 5-44 The modifyBEP command

```
modifyBEP 1 –KLM1 3 –BEP false
The cartridge encryption policy was updated successfully
```

modifyFibreChannelSettings

Use the **modifyFibreChannelSettings** command to set the ports' speed and topology for a specified drive. The speed and topology are the same for both ports in the drive.

The command uses the following syntax:

```
modifyFibreChannelSettings speed,topology -f# -c# -r#
```

The following values are valid for **speed**:

- 1
- 2
- 4
- 8
- 16 (for TS1160)

The following values are valid for **topology**:

- Auto-L
- L
- N
- Auto-N

Note: TS1160 does not support Loop mode.

The number sign (#) variables in the syntax have the following values:

-f# Frame number of the drive that contains ports 1 and 2
-c# Column number of the drive that contains ports 1 and 2
-r# Row number of the drive that contains ports 1 and 2

In Example 5-45, the **modifyFibreChannelSettings** operation sets the port speeds to 4 and the topology to N for both ports in the drive at frame 1, column 1, row 1.

Example 5-45 The modifyFibreChannelSettings command

```
modifyFibreChannelSettings 4,N -f1 -c1 -r1
The Fibre Channel settings were updated successfully
```

modifyVolserRanges

Use the **modifyVolserRanges** command to modify VOLSER ranges.

The command uses the following syntax:

```
modifyVolserRanges filename.txt
```

The **modifyVolserRanges** command receives a text file with one or more VOLSER ranges as input. Each line of the input file is interpreted as a VOLSER range to modify.

The input file is in the following format:

```
startRange,endRange,LLName,mediaType,Enable
```

The following parameter information is required:

- **startRange** is the starting VOLSER in the range to modify.
- **endRange** is the ending VOLSER in the range to modify.
- **LLName** is the name of the logical library that contains the VOLSER to modify.
- **mediaType** is `LTO` or `3592`.
- **Enable**:
  - **TRUE** = Enable this range for new cartridges only.
  - **FALSE** = Do not enable this range for new cartridges only.

The **startRange** parameter and the **endRange** parameter are used to get the range index in the logical library table. These values cannot be modified.

Example 5-46 shows an input file that instructs the system to modify the VOLSER that starts with `UAA9RH` and ends with `UBB9RH` in logical library `1` for the 3592 frame. As a result of the `TRUE` flag for the **Enable** parameter, the VOLSER range is enabled for new cartridges.

Example 5-46 The volser_ranges.txt input file

```
UAA9RH,UBB9RH,Library 1,3592,TRUE
```

Example 5-47 shows an example of the **modifyVolserRanges** command. Although each line of the input file runs successfully and each VOLSER range is created, the system returns the results to your display.

Example 5-47 The modifyVolserRanges command

```
modifyVolserRanges volser_ranges.txt
The Volser Range was modified successfully
Done
```

modifyVolserReporting

Use the **modifyVolserReporting** command to set whether the VOLSER is reported to the host in a six-character format or an eight-character format.

The command uses the following syntax:

```
modifyVolserReporting logical library name, flag value [6/8]
```

When this command runs successfully, the system returns the results of the **modifyVolserReporting** operation to your display. As shown in Example 5-48, the **modifyVolserReporting** operation instructs the system to modify the VOLSER reporting to the host for the cartridges in `Library1` to a six-character format.

Example 5-48 The modifyVolserReporting command

```
modifyVolserReporting Library1,6
Volser Reporting flag was updated successfully
```

moveFromAllDrives

Use the **moveFromAllDrives** command to eject all cartridges from all full drives and move them to the home storage location. If the home storage location is unavailable, the **moveFromAllDrives** command moves the cartridge to the first empty slot. (If you want to eject a cartridge from a specific drive, use the **moveFromDrive** command.)

The command uses the following syntax:

```
moveFromAllDrives
```

This command does not require any parameters. It performs the eject function for all of the cartridges in all of the drives where cartridges are mounted.

When the **moveFromAllDrives** command runs successfully, the system returns output to your display, as shown in Example 5-49.

Example 5-49 The moveFromAllDrives command

```
moveFromAllDrives
The drive [location] is empty now. Moving to next drive...
...
Done. All drives are empty.
```

moveFromDrive

Use the **moveFromDrive** command to eject a single cartridge from a drive and move it to the home storage location. If the home storage location is unavailable, the **moveFromDrive** command moves the cartridge to the first empty slot. (If you want to eject a cartridge from the drive, and move it to a specific location, use the **moveToDrive** command.)

The command uses the following syntax:

```
moveFromDrive -f# -c# -r#
```

The number sign (#) variables in the syntax specify the drive where the cartridge is located before you eject it. The variables have the following values:

-f# Frame number of the moveFrom location
-c# Column number of the moveFrom location
-r# Row number of the moveFrom location

When the **moveFromDrive** command successfully runs and the cartridge is successfully unmounted from the drive, the system returns output to your display. Example 5-50 shows that the system is instructed to eject the cartridge from the drive in frame `1`, column `4`, row `1`, and place it in a storage slot.

Example 5-50 The moveFromDrive command

```
moveFromDrive -f1 -c4 -r1
The cartridge was ejected successfully
```

moveToDrive

Use the **moveToDrive** command to move a cartridge from a specified location to another specified location.

The command uses the following syntax:

```
moveToDrive F#,C#,R#,T# -f# -c# -r#
```

You must first specify the drive location of the cartridge that you want moved. Then, specify the destination where you want to move the cartridge.

The number sign (#) variables in the syntax have the following values:

F# Frame number of the drive for the moveFrom location
C# Column number of the drive for the moveFrom location
R# Row number of the drive for the moveFrom location
T# Tier number of the drive for the moveFrom location
-f# Frame number of the drive for the moveTo location
-c# Column number of the drive for the moveTo location
-r# Row number of the drive for the moveTo location

> **Note:** The CLI is not case-sensitive. In the **moveToDrive** command, uppercase and lowercase letters are used to distinguish between the moveFrom location and the moveTo location.

When the **moveToDrive** command successfully runs and the cartridge is successfully mounted in the new location, the system returns output to your display.

Example 5-51 shows the command that instructs the system to move the cartridge in drive `F1`, `C5`, `R1`, `T1` to the drive in `f1`, `c4`, `r1`. If the cartridge destination is in the same tier as the moveFrom location, you do not need to specify the tier number in the moveTo location.

Example 5-51 The moveToDrive command

```
moveToDrive F1,C5,R1,T1 -f1 -c4 -r1
The cartridge was mounted successfully
```

prestageDataCartridges

Use the **prestageDataCartridges** command to move cartridges from the high-density slots to cartridge cache locations, as specified in a text file or by using a list of VOLSERs. All data cartridges that are being prestaged must be in tier 3 or higher.

The command uses the following syntax:

```
prestageDataCartridges filename.txt or
prestageDataCartridges [VOLSER1],[VOLSER2],[VOLSERX],...
```

The **prestageDataCartridges** command can use a text file with one or more prestage operations as input. The `filename` variable specifies the file name of the input file.

The text file is in one of the following formats:

- From a storage location: `[F,C,R,T]`
- VOLSER: `[VOLSER]`

Each line in the text file is interpreted as a prestage operation. In the text file, the data cartridge location can be specified with the `[F,C,R,T]` location or the `VOLSER`.

> **Note:** All cartridges that are being prestaged must be in tier 3 or higher.

Example 5-52 shows a **prestageDataCartridges** input file with the `F,C,R,T` reference.

Example 5-52 The prestageDataCartridges input file with the F,C,R,T reference

```
F1,C2,R8,T3
F1,C3,R18,T5
```

Example 5-53 shows the **prestageDataCartridges** input file with the `VOLSER` reference.

Example 5-53 The prestageDataCartridges input file with the VOLSER reference

```
UAA9RHL7
UBB9RHL7
```

When the **prestageDataCartridges** command runs successfully, the system returns output to your display.

Example 5-54 shows the command that instructs the system to prestage the data cartridges to the correct library location that is based on the information that is specified in the `prestage.txt` file.

Example 5-54 The prestageDataCartridges command

```
prestageDataCartridges prestage.txt
The cartridge [location] was moved successfully. Next element...
...
Done
```

removeDataCartridges

Use the **removeDataCartridges** command to remove the data cartridges from a storage location to the I/O station. If VIO is enabled for a specific cartridge, the **removeDataCartridges** command does not move the cartridge to the I/O station. Instead, the cartridge is marked as a cartridge that is ready for export. The cartridge appears as unassigned. If VIO is disabled, the cartridge moves to the I/O station.

The command uses the following syntax:

```
removeDataCartridges filename.txt
```

The **removeDataCartridges** command receives a text file with one or more remove operations as input. The text file is in the following format:

- `[F,C,R,T]`
- `VOLSER`

Each line in the text file is interpreted as a remove operation.

Example 5-55 shows a **removeDataCartridges** input file. In the first line of the text file, the data cartridge in storage slot location frame `1`, column `2`, row `8`, tier `0` is moved to the I/O station. In the second line of the text file, the data cartridge in storage slot location frame `1`, column `3`, row `18`, tier `5` is moved to the I/O station.

Example 5-55 The removeDataCartridges input file

```
F1,C2,R8,T0
F1,C3,R18,T5
UBB9RHL7
```

Example 5-56 shows the command that instructs the system to remove the data cartridges from a storage slot to the I/O station based on the information in the `todays_removes.txt` file. Example 5-56 shows the **`removeDataCartridges`** output.

Example 5-56 The removeDataCartridges output

```
removeDataCartridges todays_removes.txt
The cartridge [FCRT] has been removed successfully
...
Done
```

resetDrive

Use the **`resetDrive`** command to power-cycle a specific drive.

The command uses the following syntax:

```
resetDrive -f# -c# -r#
```

The number sign (#) variables specify the location of the drive and feature the following values:

`-f#` Frame number of the drive
`-c#` Column number of the drive
`-r#` Row number of the drive

When the **`resetDrive`** command runs successfully, the system returns output to your display, as shown in Example 5-57.

Example 5-57 The resetDrive command output

```
resetDrive -f1 -c4 -r1
The drive was reset
successfully
```

resetNodeCards

Use the **`resetNodeCards`** command to reset one or more specified node cards.

> **Note:** No warning to cancel is shown; therefore, use this command with care because running this command immediately resets the selected node card.

The command uses the following syntax:

```
resetNodeCards node_name_1, node_name_2, . . .
```

The parameters for this command are items in a list of comma-separated names of nodes to reset.

The following list shows the possible values of node names:

- `ALL`
- `ALLLCA`
- `ALLACC`
- `ALLSMC`
- `ALLXYC`
- `ALLPLUS`
- `XYCPLUS`
- `ACCPLUS`

- `LCAPLUS`
- `XYA`
- `XYB`
- `OPC1`
- `SMC[2-16]`
- `ACCA`
- `ACCB`
- `LCA[1-31]`
- `LCAB[1-31]`

When the **`resetNodeCards`** command runs successfully, the system returns output to your display. Example 5-58 shows the command where the system is instructed to reset the node cards in nodes `LCA1`, `LCAB3`, and `ACCA`.

Example 5-58 The resetNodeCards command

```
resetNodeCards LCA1,LCAB3,ACCA
The command was executed successfully
```

restoreConfiguration

Use the **`restoreConfiguration`** command to restore the configuration database to the library. The configuration file (`configuration.db`) is uploaded to the library.

The command uses the following syntax:

```
restoreConfiguration the_name_of_the configuration_file.db
```

You must specify the name of a configuration file to upload. The file name can be any name that you choose, but the file must have a `.db` extension. If the file is stored in the same directory where you are running the CLI, you do not need to enter a path name for the file. However, if the file is stored in a different directory, it is necessary to enter a path name for the file.

When the **`restoreConfiguration`** command runs successfully, the system returns the result of the operation to your display, as shown in Example 5-59.

Example 5-59 The restoreConfiguration command

```
restoreConfiguration myconfig.db
Name: myconfig.db
>>>Uploading file...
Done
..
The configuration database has been restored successfully.
```

saveConfiguration

Use the **`saveConfiguration`** command to save the configuration database to the user's computer. The configuration file (`configuration_backup.db`) is saved in the same directory in which the CLI is running. This command requires a service user and password.

The command uses the following syntax:

```
saveConfiguration
```

No parameters are necessary. When the **`saveConfiguration`** command runs successfully, the system returns the result of the operation to your display, as shown in Example 5-60 on page 359.

Example 5-60 The saveConfiguration command

```
saveConfiguration
Downloading... [completed] / [totalSize] has been downloaded.
..
The configuration file: configuration_backup.db has been downloaded.
```

setAccessorZones

This command is for use on Dual accessor libraries on and used to specify which tape drives and I/O stations (within sets of frames) are served by which accessor. This command can also be used to deactivate an accessor.

By default, each accessor is zoned to serve half of the frames in the library. Therefore, in a 12-frame library, Accessor A serves frames 1 - 6 and Accessor B serves frames 7 - 12.

You can view the current accessor zone settings by using the **viewAccessorZones** command. You can also view these settings in the TS4500 management GUI in the **Settings** → **Library** → **Cartridges and Accessors** page. The Accessor Preferred Zones graphic is refreshed to reflect changes made both through the GUI and with the setAccessorZones command.

Elastic capacity settings that are made through the Cartridges and Accessors page are unaffected by zone changes that are made with the **setAccessorZones** command.

Use the **no** parameter to specify a different zone setting for an accessor, or to deactivate an accessor. The use of this command to deactivate an accessor is the same as deactivating it in the Cartridges and Accessors GUI page.

The command uses the following syntax:

```
setAccessorZones no,[frameNumber|0|255]
```

The parameters in the syntax have the following values:

- **no**: Do not use the default zone setting. Instead, use the setting specified here.
- **frame Number from 1 to 18**: Defines the last frame served by Accessor A. All remaining frames in the library are served by Accessor B:
  - **0**: Deactivate Accessor A.
  - **255**: Deactivate Accessor B.

The example that is shown in Example 5-61 shows restricting Accessor A to frames 1, 2, 3, 4, and 5. Accessor B serves the rest of the frames in the library.

Example 5-61 setAccessorZones to frame 5

```
setAccessorZones No,5
The accessor zones was updated successfully
```

The example that is shown in Example 5-62 shows restricting Accessor A to frames 1, 2, 3, 4, and 5. Accessor B serves the rest of the frames in the library.

Example 5-62 setAccessorZones default

```
setAccessorZones Yes,255
The accessor zones was updated successfully
```

setAutoEjectCleaningCarts

Use the **`setAutoEjectCleaningCarts`** command to set the `enabled` or `disabled` flag for the auto eject expired cleaning cartridges function in the library.

The command uses the following syntax:

```
setAutoEjectCleaningCarts enabled or disabled
```

Set the auto eject expired cleaning cartridges function flag:

- `enabled` = on
- `disabled` = off

When the **`setAutoEjectCleaningCarts`** command runs successfully, the system returns output to your display, as shown in Example 5-63.

Example 5-63 The setAutoEjectCleaningCarts command

```
setAutoEjectCleaningCarts enabled
The auto eject cleaning cartridges flag was set successfully
```

setDrivePortsId

Use the **`setDrivePortsId`** command to set the ports' IDs for a specified drive (Port 1 and Port 2).

The command uses the following syntax:

```
setDrivePortsIdPort1ID,Port2ID -f# -c# -r#
```

The number sign (#) variables in the syntax have the following values:

-f#	Frame number of the drive that contains ports 1 and 2.
-c#	Column number of the drive that contains ports 1 and 2.
-r#	Row number of the drive that contains ports 1 and 2.

When the **`setDrivePortsId`** command runs successfully, the system returns the result of the operation to your display, as shown in Example 5-64.

Example 5-64 The setDrivePortsId command

```
setDrivePortsId 18,21 -f1 -c1 -r1
The ports IDs were updated successfully
```

setDriveUse

Use the **`setDriveuse`** command to set and configure the drive for media access, control path with media access, or media verification.

The command uses the following syntax:

```
setDriveUse -use [useaccess, controlPath, verification] -f# -c# -r#
```

The number sign (#) variables in the syntax have the following values:

-f#	Frame number of the drive.
-c#	Column number of the drive.
-r#	Row number of the drive.

When the **`setDriveUse`** command runs successfully, the system returns the result of the operation to the display, as shown in Example 5-65 on page 361.

Example 5-65 The setDrivePortsId command

```
--setDriveUse -use controlPath -f1 -C4 -R1
The drive was successfully updated
```

setISCSI

The **setISCSI** command updates the network configuration settings for a single Ethernet drive and to optionally change the name and alias of the drive. Each port can be configured separately by using this command, whereas network changes that are made through the TS4500 management GUI are always applied to both ports.

To look up a drive's current settings, use the **viewISCSI** command or refer to the **Drives** → **Ethernet Ports** GUI page

The command uses the following syntax:

```
setISCSI -portxAddress DHCP|[ip-address] -subnet [subnet-mask] -gateway
[gateway-address] -name DEFAULT|[iSCSI-name] -alias [iSCSI-alias] -f# -c# -r#
```

The parameters have the following values:

- **portxAddress DHCP | [ip-address]**

  Both ports can be modified at the same time; portxAddress is **port0Address** for port 0 or **port1Address** for port 1.

  DHCP enables DHCP for the port selected.

- **[ip-address]** disables DHCP and assigns this value as the IP address for the specified port. (IPv4 format). You can specify the settings for one port or both ports in the same command.
- **subnet [subnet-mask]**. This is required if DHCP is being disabled. This is the subnet mask for the specified port.
- **gateway [gateway-address]**. This value is required if DHCP is being disabled. It is the gateway for the specified port.
- **name DEFAULT|[iSCSI-name]**. This value is optional. It is the name for the drive.

  DEFAULT restores the drive name to the system-assigned name.

  SCSI-name is the new name for the drive. A port ID is required when specifying this parameter, even though the name is applied to the drive. You can specify either port.

- **alias [iSCSI-alias]** is an optional parameter and sets a new alias for the drive. A port ID is required when specifying this parameter, even though the alias is applied to the drive. You can specify either port.
- **[-f# -c# -r#]** Is a required field and is the location of the drive, where:
  - **-f#** = frame number
  - **-c#** = column number
  - **-r#** = row number

As shown in Example 5-66, the use of the command enables DHCP for port 0 of the drive in F1C4R1 and renames the drive to new-drive-name.

Example 5-66 setISCSI default command

```
setISCSI -port0Address DHCP -name new-drive-name -f1 -c4 -r1
The iSCSI Drive Port Settings were updated.
```

As shown in Example 5-67, the use of the command disables DHCP for port 0 of the drive in F1C4R1. It also assigns the address 9.1.2.4 to port 0 and also assigns the provided subnet mask and gateway to the port.

Example 5-67 setISCSI set IP command

```
setISCSI -port0Address 9.1.2.4 -subnet 255.255.255.0 -gateway 9.1.0.100 -f1 -c4
-r1
The iSCSI Drive Port Settings were updated.
```

As shown in Example 5-68, the use of the command enables DHCP for port 0 and disables DHCP for port 1 of the drive in F1C4R1, and changes the alias for the drive to new-drive-alias.

Example 5-68 setISCSI alias command

```
setISCSI -port0Address DHCP -port1Address 9.1.2.5 -subnet 255.255.255.1 -gateway
9.1.0.100 -alias new-drive-alias -f1 -c4 -r1
The iSCSI Drive Port Settings were updated.
```

setLibraryTime

The **setLibraryTime** command takes the date, time, and time zone in the server where the CLI is running and sets the library with those parameters.

The command uses the following syntax:

```
setLibraryTime
```

No parameters are necessary. When the **setLibraryTime** command runs successfully, the system returns the result of the operation to your display, as shown in Example 5-69.

Example 5-69 The setLibraryTime command

```
setLibraryTime
The library date and time were updated successfully
```

setMacAddress

The **setMacAddress** command changes the MAC address for an Ethernet port (or both ports). To look up a port's current MAC address, use the **viewMacAddress** command or refer to the **Drives** → **Ethernet Ports** GUI page.

The command uses the following syntax:

```
setMacAddress -portxAddress DEFAULT|[MAC-address]} -portxAddress
DEFAULT|[MAC-address] <-f# -c# -r#>
```

The parameters have the following values:

- **portxAddres**s, Both ports can be modified at the same time. portxAddress is **port0Address** for port 0 or **port1Address** for port 1.
- **DEFAULT** returns the MAC address to its default setting.
- [MAC-address] is the new MAC address, expressed as a 12-digit hexadecimal number. It can be entered either with or without colons for example, 00:1a:64:eb:04:83 and 001a64eb0483 are both acceptable.

- **-f# -c# -r#** is required and is the location of the drive:
  - **-f#** is the frame number
  - **-c#** is the column number
  - **-r#** is the row number

When the **setMacAddress** command runs successfully, the system returns the result of the operation to your display, as shown in Example 5-70. This example shows both of the accepted formats of the MAC address.

Example 5-70 The setMacAddress command

```
setMacAddress -port0Address 00:1a:64:eb:04:83 -port1Address 001a64eb0484 -f1 -c4
-r1
MAC address(es) updated.
```

setMaximumVIOCartridges

Use the **setMaximumVIOCartridges** command to update the number of virtual I/O slots for the logical library.

The command uses the following syntax:

```
setMaximumVIOCartridges logical library name, value [32-255]
```

The maximum number of cartridges can be any value 32 - 255. When this command runs successfully, the system returns the results of the **setMaximumVIOCartridges** operation to your display.

Example 5-71 shows the command to set the maximum virtual I/O cartridges in `Library1` to `200`.

Example 5-71 The setMaximumVIOCartridges command

```
setMaximumVIOCartridges Library1,200
The MAX VIO was updated successfully
```

setNMADetection

Use the **setNMADetection** command to set (on or off) the No Motion Allowed (NMA) detection flag in the library.

The command uses the following syntax:

```
setNMADetection TRUE or FALSE
```

The parameter can have either of the following values:

`TRUE` NMA detection is on.
`FALSE` NMA detection is off.

When the **setNMADetection** command runs successfully, the system returns the result of the operation to your display, as shown in Example 5-72.

Example 5-72 The setNMADetection command

```
setNMADetection TRUE
The NMA detection flag was updated successfully
```

setRolePermissions

Use the **setRolePermissions** command with the output from the **viewRolePermissions** command. This command instructs the system to change the permissions for a specified role.

The command uses the following syntax:

```
setRolePermissions filename_of_role_permissions.txt -role role name
```

The **setRolePermissions** command receives a text file (which was created by the **viewRolepermissions** command) as input. Each line in the text file is interpreted as an instruction to set a permission. The text file can be edited to change any of the access levels for the listed pages.

The following access levels are valid:

Read only The user has read-only access to the web page.
Modify The user can run the edit operation for the web page.
No Access The user cannot view the web page.

You can edit the text file to contain only the names of the web pages and permissions to change.

Example 5-73 shows the input file that was originally generated by the **viewRolePermissions** command. In this case, the input file is unchanged. However, you can edit the input file to contain only the specific permissions to change with the **setRolePermissions** command.

Example 5-73 The setRolePermissions input file

```
                           Action,   Access Level
                              AIT,      No Access
                       Cartridges,         Modify
    Cartridges by Logical Library,         Modify
            Cleaning Cartridges,          Modify
                  Debug Messages,      No Access
                           Drives,         Modify
        Drives by Logical Library,         Modify
             Email Notifications,      Read Only
                Email Recipients,      Read Only
             Encryption Internal,      Read Only
          Encryption Key Manager,         Modify
                  Ethernet Ports,         Modify
                          Events,         Modify
             Library Information,         Modify
              Licensed Functions,         Modify
               Logical Libraries,         Modify
         Management GUI Behavior,         Modify
                  Master Console,      No Access
                  Password Rules,         Modify
                Position Control,      No Access
           Remote Authentication,         Modify
                           Roles,         Modify
               SNMP Destinations,         Modify
              SNMP Notifications,         Modify
                      Scan Speed,      No Access
             Secure Socket Layer,         Modify
                    Service Port,      No Access
            Syslogs Notifications,         Modify
                          System,         Modify
```

```
                     System,         Modify
       System Date and Time,         Modify
                      Tasks,         Modify
                      Users,         Modify
              VOLSER Ranges,         Modify
VOLSER Ranges by Logical Library,    Modify
```

When the **setRolePermissions** command runs successfully, the system returns output to your display. Example 5-74 shows the command that instructs the system to set the role permissions for the administrator role, according to information in the `administrator_permissions.txt` file.

Example 5-74 The setRolePermissions command

```
setRolePermissions administrator_permissions.txt role administrator
The permissions were updated successfully
```

setScannerSpeed

Use the **setScannerSpeed** command to set the speed of the scanner.

The command uses the following syntax:

```
setScannerSpeed speed
```

The following values for the `speed` variable are valid:

- `0` (nominal speed)
- `30`
- `40`
- `50`
- `60`
- `70`
- `80`
- `90`
- `100`

When the **setScannerSpeed** command runs successfully, the system returns output to your display. Example 5-75 shows the command that instructs the system to set the scanner speed to `100`.

Example 5-75 The setScannerSpeed command

```
setScannerSpeed 100
The scanner speed was updated successfully
```

setSlotOffline

Use the **setSlotOffline** command to set a storage slot offline.

The command uses the following syntax:

```
setSlotOffline -f# -c# -r#
```

The number sign (#) variables are defined in the following manner:

-f# Frame number of the slot
-c# Column number of the slot
-r# Row number of the slot

Note: When you use `r` =0, all rows for the specific column are set offline.

When the **`setSlotOffline`** command runs successfully, the system returns output to your display. Example 5-76 shows that the command was successful.

Example 5-76 The setSlotOffline command

```
setSlotOffline -f1 -c4 -r3
Storage Slot was set offline
```

setSlotOnline

Use the **`setSlotOnline`** command to set a storage slot online.

The command uses the following syntax:

```
setSlotOnline -f# -c# -r#
```

The number sign (#) variables are defined in the following manner:

`-f#` Frame number of the slot
`-c#` Column number of the slot
`-r#` Row number of the slot

Note: When you use **`r =0`**, all rows for the specific column are set online.

When the **`setSlotOnline`** command runs successfully, the system returns output to your display. Example 5-77 shows that the command was successful.

Example 5-77 The setSlotOnline command

```
setSlotOnline -f1 -c4 -r3
Storage Slot was set online
```

setSSL

Use the **`setSSL`** command to set the Secure Sockets Layer (SSL) flag to `enabled` or `disabled`.

The command uses the following syntax:

```
setSSL enabled or disabled
```

When the **`setSSL`** command runs successfully, the system returns output to your display. Example 5-78 shows the command that instructs the system to set the SSL flag to `enabled`.

Example 5-78 The setSSL command

```
setSSL enabled
The SSL flag was updated successfully
```

setUtilThreshold

Use the **`setUtilThreshold`** command to change the default values of the Capacity Utilization Threshold and the Dual Accessor Utilization Threshold.

If the number of cartridges exceeds the defined capacity utilization threshold, the Capacity POD in the monitor system page turns yellow. If this scenario occurs, it is possible to remove cartridges, increase capacity threshold, or purchase more capacity. By using this command, you can increase or decrease the value where the alert occurs.

The dual accessor utilization threshold (`-daut`) option is applicable to only the dual accessor libraries with elastic capacity enabled. It determines the threshold at which cartridges are moved to accessor service areas in the end frames. The default for DAUT is 98% and CUT is 99%.

The command uses the following syntax:

```
setUtilThreshold -daut # -cut #
```

The `-cut`, `-daut`, or both options can be used. Consider the following points:

- `-daut` values must be 90 - 100. Decimal places are not supported.
- The following `-cut` values can be used:
  - 50, 55, 60, 65, 70, 75, 80, 85, and 90
  - 91, 92, 93, 94, 95, 96, 97, and 98
  - 99.0, 99.1, 99.2, 99.3, 99.4, 99.5, 99.6, 99.7, 99.8, and 99.9

When the **`setUtilThreshold`** command runs successfully, the system returns output to the display, as shown in Example 5-79.

Example 5-79 The setUtilThreshold command

```
setUtilThreshold -cut 97 -daut 97
```

Use the **`viewUtilityThreshold`** command to display current settings.

showQueuedExports

Use the **`showQueuedExports`** command to set a flag to show or hide the list of queued exports. This command is useful to view the status of the recent export jobs, cancel active jobs, or remove any recent jobs from the list.

The command uses the following syntax:

```
showQueuedExports logical library name, flag value [true / false]
```

A flag value of `true` *shows* the queued exports list. A flag value of `false` *hides* the queued exports list.

When this command runs successfully, the system returns the results of the **`showQueuedExports`** operation to your display. Example 5-80 shows the command that shows the queued exports list for the `Library1` logical library.

Example 5-80 The showQueuedExports command

```
showQueuedExports Library1, true
Queued exports flag was updated successfully
```

startInventory

Use the **startInventory** command to start the library inventory process.

The command uses the following syntax:

```
startInventory --library SINGLE|NO --audit FIRST|ALL> --frame ALL|<frameIdx>
```

To audit only tier 0 and 1, select `--audit FIRST` parameter. Because the `--audit ALL` parameter performs an inventory with full audit, take care when this parameter is used because as it can take up to 45 minutes per frame to complete.

When the **startInventory** command runs successfully, the system returns output to the display, as shown in Example 5-81.

Example 5-81 The startInventory command

```
startInventory --library SINGLE --audit FIRST --frame 1
Inventory has started
```

startCalibration

Use the **startCalibration** command to start calibration of the library, or components inside the library.

The command uses the following syntax:

```
startCalibration FRAME,<frame number>
                 DRIVE, <f#r#c#>
                 IO, <frameIndex,IOIndex>
                 ACCESSOR,<1=A,2=B>
```

When the **startCalibration** command runs successfully, the system returns output to the display, as shown in Example 5-82.

Example 5-82 The startCalibration command

```
startCalibration DRIVE,f2c3r1
Calibration has started
```

startDriveService

Use the **startDriveService** command to prepare a drive service action. The drive is not usable for applications. The drive location mist be specified. A task is created and Progress can be monitored at the GUI.

The command uses the following syntax:

```
startDriveService -f# -c# -r#
```

The number sign (#) variables specify the location of the drive, and have the following values:

- **-f#** Frame number of the drive
- **-c#** Column number of the drive
- **-r#** Row number of the drive

When the **startDriveService** command runs successfully, the system returns output to your display, as shown in Example 5-83.

Example 5-83 The startDriveService command

```
startDriveService -f3 -c2 -r1
The drive Service has started
```

startDiscoverHW

Use the **startDiscoverHW** command to start the library hardware discovery. The following options are available:

- `OVERWRITE <frames>` deletes all existing configuration data for the specified frames and re-creates the data.
- `<frames>` saves the existing configuration data for the specified frames and adds information for any new hardware that is discovered.

The command uses the following syntax:

```
startDiscoverHW OVERWRITE,f#,f#,f#
```

or

```
startDiscoverHW f#,f#,f#
```

When the **startDiscoverHW** command runs successfully, the system returns output to the display, as shown in Example 5-84.

Example 5-84 The startDiscoverHW command

```
startDiscoverHW 1,2,3
Discover HW has started
```

startLibraryVerify

Use the **startLibraryVerify** command to start the library verification process for I/O stations. When the command completes, you are prompted to open all I/O station doors and then run the **continueLibraryVerify** command. Library verification is a three-step process in which the following commands must be run in the order in which they are listed:

- `startLibraryVerify`
- `continueLibraryVerify`
- `continueCloseLibraryVerify`

The command uses the following syntax:

```
startLibraryVerify
```

When the **startLibraryVerify** command runs successfully, the system returns output to the display, as shown in Example 5-85.

Example 5-85 The startLibraryVerify command

```
startLibraryVerify
Library verify - IO stations - has started...Library verify in progress
..........................................................................
..........Done!
Next step: OPEN all IO doors , and then use "continueLibraryVerify" command to
continue the test
```

unassignDrive

Use the **unassignDrive** command to unassign a specific drive.

The command uses the following syntax:

```
unassignDrive F#C#R#
```

The number sign (#) variables are defined in the following manner:

-f# Frame number of the drive
-c# Column number of the drive
-r# Row number of the drive

When the **unassignDrive** command runs successfully, the system returns output to the display, as shown in Example 5-86.

Example 5-86 The unassignDrive command

```
unassignDrive F2C3R4
The drive was unassigned successfully
```

version

Use the **version** command to view the current version of the Gen 4 TS4500 Web Interface CLI program.

The command uses the following syntax:

```
version
```

When the **version** command runs successfully, the system returns output to your display, as shown in Example 5-87.

Example 5-87 The version command

```
version
CLI GEN 4 Version: 1.8.0.1
Build: 07/20/2021
```

viewAccessor

Use the **viewAccessor** command to show the accessor's status and usage statistics. This command features no parameters.

When the **viewAccessor** command runs successfully, the system returns output to your display, as shown in Example 5-88.

Example 5-88 The viewAccessor command

```
viewAccessor
Availability
                   ,     Accessor A
           Accessor,    OK - Online
          Gripper 1,    OK - Online
          Gripper 2,    OK - Online

Usage Statistics
                        ,     Accessor A Component
                  Pivots,                      195
```

```
     Gripper 1 gets,                        5
     Gripper 1 puts,                        5
     Gripper 2 gets,                        1
     Gripper 2 puts,                        1
     Bar code scans,                     9355
  X travel (meters),                      112
  Y travel (meters),                      138
```

viewAccessorZones

Use the **viewAccessorZones** command to view which tape drives and I/O stations (within sets of frames) are served by which accessor. You can also view these settings in the TS4500 management GUI in the **Settings** → **Library** → **Cartridges and Accessors** page.

The command uses the following syntax:

```
viewAccessorZones > outputfilename.csv
```

Use the **viewAccessorZones > outputilename.csv** command to instruct the system to generate a .csv file from the output by entering a name as a parameter to the command for the file that is generated, as shown in the following example:

viewNodeCards > *any_name_you_want*.csv

The system generates the .csv file and stores it in the directory that contains the TS4500CLI.jar file. No message is returned to your display. To store the .csv file in another directory, specify the name of the directory in the command, as shown in Example 5-89.

Example 5-89 The viewAccessorZones command

```
viewAccessorZones
Accessor A: 1 - 12
Accessor B: 13 - 16
```

viewAdvancedEncryptionSettings

Use the **viewAdvancedEncryptionSettings** command to view high-level information for the encryption settings for a specific logical library.

The command uses the syntax:

```
viewAdvancedEncryptionSettings [Logical Library Name]
```

The **viewAdvancedEncryptionSettings** command displays information for all of the encryption settings for a specific logical library that is named Redbooks, as shown in Example 5-90.

Example 5-90 The viewAdvancedEncryptionSettings command

```
viewAdvancedEncryptionSettings Redbooks
              Name:                             Redbooks
   Advanced Method:                No advanced setting
   Advanced Policy:                No advanced setting
      Density Code:                No advanced setting
          Key Path:                No advanced setting
```

viewBEP

Use the **viewBEP** command to show a list of all VOLSER ranges that shows the bar code encryption policy (BEP) settings. This command features no parameters.

When the **viewBEP** command runs successfully, the system returns output to your display, as shown in Example 5-91.

Example 5-91 The viewBEP command

```
viewBEP
Index,Volser Start, Volser End,  Logical Library, Media Type,  KLM1, KLM2,  BEP
    1,      JK0000,     JK0009,         SLS_JAG1,       3592,     0,    0,    0
    2,      RBD000,     RBD099,  Redbooks_custom,       3592,     0,    0,    0
    3,      RBA000,     RBA999,    LTFS_redbooks,       3592,     0,    0,    0
    4,      DF0000,     DF9999,  Redbooks_custom,       3592,     0,    0,    0
```

viewCleaningCartridges

Use the **viewCleaningCartridges** command to view high-level information for all cleaning cartridges in a library.

The command uses the following syntax:

```
viewCleaningCartridges
```

The **viewCleaningCartridges** command does not require any parameters. It instructs the system to display information for all of the cleaning cartridges in the library. When the **viewCleaningCartridges** command runs successfully, the system returns output to your display.

The output file sorts the names of the cleaning cartridges by VOLSER, as shown in Example 5-92.

Example 5-92 The viewCleaningCartridges command

```
viewCleaningCartridges
Volume Serial, Logical Library,Element Address, Media Type, Location(F,C,R),Cleans
remaining, Most Recent use
 CLN305L1,       255,      null,                  LTO,     Slot(F1,C6,R16,T0),
0,         04/25/2014 00:50:04
 CLN510L1,       255,      null,                   LTO,     Slot(F1,C6,R12,T0),
25,         05/08/2014 05:35:54
```

The same command can generate a `.csv` file that is downloaded into the directory of the `TS4500CLI.jar` file. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

```
viewCleaningCartridges > any_name_you_want.csv
```

No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command. The command is shown in Example 5-93.

Example 5-93 Use viewCleaningCartridges to generate a .csv file

```
viewCleaningCartridges > C:\CLI\CVSFiles\cleaning_cartridges.csv
```

Note: All view commands can generate a `.csv` file by adding the **> *any_name_you_want*.csv** parameter to the end of the command.

viewDataCartridges

Use the **viewDataCartridges** command to view the high-level information for all data cartridges in a library.

The command uses the following syntax:

```
viewDataCartridges
```

The **viewDataCartridges** command does not require any parameters. It lists the information for all data cartridges in the library. When the **viewDataCartridges** command runs successfully, the system returns output to your display. The output file sorts the names of the data cartridges by VOLSER, as shown in Example 5-94.

Example 5-94 The viewDataCartridges command

```
viewDataCartridges
Volume Serial, Logical Library,Element Address, Media Type,
Location(F,C,R),Encryption, Most Recent use
IM1338L7,      Unassigned,      Unknown,           LTO, Slot(F1,C3,R1,T2),
Unknown,       04/24/2014 08:39:34
IM1364L7,      FC_LTO6_5856,          1054,        LTO,Slot(F4,C8,R38,T0), Not
Encrypted, 05/07/2014 15:24:24
IM1364L7,      FC_LTO6_5856,          1265,        LTO,Slot(F4,C8,R38,T0), Not
Encrypted, 05/07/2014 15:24:24
```

You can use the **viewDataCartridges** command to instruct the system to generate a `.csv` file from the output by entering a name as a parameter to the command for the file that is generated.

The command uses the following syntax:

```
viewDataCartridges > any_name_you_want.csv
```

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-95.

Example 5-95 The viewDataCartridges command with a .csv file

```
viewDataCartridges > C:\CLI\CVSFiles\todays_data_cartridges.csv
```

> **Note:** All view commands can generate a `.csv` file by adding the **> *any_name_you_want*.csv** parameter at the end of the command.

viewDriveDetails

The **viewDriveDetails** command provides a display of key information about the drive that is useful for troubleshooting. It displays the contents of the drive, name of the logical library, worldwide node numbers (WWNNs), port ID information, and so on.

The command uses the following syntax:

```
viewDriveDetails -f# -c# -r#
```

The number sign (#) variables specify the location of the drive, and have the following values:

- **-f#** Frame number of the drive
- **-c#** Column number of the drive
- **-r#** Row number of the drive

When the **viewDriveDetails** command runs successfully, the system returns output to your display, as shown in Example 5-96.

Example 5-96 The viewDriveDetails command

```
viewDriveDetails -f 1 -c 4 -r  1
 Location(F,C,R)        F1, C4, R1
           State            ONLINE
            Type           3592-E07
 Logical Library     Redbooks_custom
    Control Path           Enabled
        Contents             Empty
        Firmware             3DCB
            WWNN    50050763001a000c
 Element Address               257
   Drive Display                ND
          Port 0
         Loop ID                13
            WWPN          005a000c
          Port 1
         Loop ID                77
            WWPN          009a000c
```

The **viewDriveDetails** command includes a parameter that instructs the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

`viewDriveDetails >` *`any_name_you_want`*`.csv`

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to the display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-97.

Example 5-97 The viewDriveDetails command with a .csv file

```
viewDriveDetails -f1 -c1 -r3 > C:\CLI\CVSFiles\drive_details.csv
```

> **Note:** All view commands generate a `.csv` by adding the **>** ***any_name_you_want*.csv** parameter at the end of the command.

viewDrivePod

Use the **viewDrivePod** command to display the drive portion of the status Pods. Pods in the bottom of the GUI System page show a quick view of capacity, drive utilization, and library health status. This pod is for drive status. The command uses the following syntax:

`viewDrivePod`

This command does not require any parameters. It lists a summary of information for all drives in the library. When the **viewDrivePod** runs successfully, the system returns output to the display, as shown in Example 5-98 on page 375.

Example 5-98 The viewDrivePod command

```
viewDrivePod

[{"clazz":"com.ibm.storage.anaconda.events.DriveStationEvent","drivesInUse":0,"lib
raryIdx":-1,"libraryName":"","podStatus":"HEALTHY","totalDrives":12,"topic":"DRIVE
_STATION","id":70800,"arguments":null,"timestamp":1538729795448,"topic":"DRIVE_STA
TION"}]
```

> **Note:** All view commands generate a `.csv` by adding the **`>`** ***`any_name_you_want`*****`.csv`** parameter at the end of the command.

viewDriveSummary

Use the **`viewDriveSummary`** command to show all of the drives in the library. This summary information is useful to network administrators who decide where to mount a cartridge. If you want to list detailed information about a specific drive, use the **`viewDriveDetails`** command.

The command uses the following syntax:

```
viewDriveSummary
```

This command does not require any parameters. It lists a summary of information for all drives in the library. When the **`viewDriveSummary`** runs successfully, the system returns output to your display, as shown in Example 5-99.

Example 5-99 The viewDriveSummary command

```
viewDriveSummary
Location(F,C,R), State,   Type, Contents, Firmware, Serial,  WWWNN,Element
Address,Logical Library
  F1, C4, R1,   ONLINE,3592-E07,  Empty,   3DCB , 0000013B008E, ,
257,Redbooks_custom
  F1, C4, R4,   ONLINE,3592-E07,  Empty,   3DCB , 0000013B006F, ,
258,Redbooks_custom
```

The **`viewDriveSummary`** command can include a parameter that instructs the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

```
viewDriveSummary > any_name_you_want.csv
```

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-100.

Example 5-100 The viewDriveSummary command with a .csv file

```
viewDriveSummary > C:\CLI\CVSFiles\drive_summary.csv
```

viewDriveVPD

Use the **viewDriveVPD** command to show vital product data (VPD) information for all drives. The **viewDriveVPD** command is useful if an administrator or IBM service support representative (SSR) needs to see information, such as a product model number, product release level, or other information that is specific to the device type.

The command uses the following syntax:

```
viewDriveVPD
```

The **viewDriveVPD** command does not require any parameters. It lists all of the vital product data for all of the drives in the library.

When the **viewDriveVPD** command runs successfully, the system returns output to your display, as shown in Example 5-101.

Example 5-101 The viewDriveVPD command

```
viewDriveVPD
Location(F,C,R), Drive Type,Firmware Version, Machine Type, Serial number
 F1, C4, R1,     3592-E07,      3DCB ,        3592,         00078B008E
 F1, C4, R4,     3592-E07,      3DCB ,        3592,         00078B0060
```

The **viewDriveVPD** command can include a parameter that instructs the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewDriveVPD > *any_name_you_want*.csv

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-102.

Example 5-102 The viewDriveVPD command with a .csv file

```
viewDriveVPD  > C:\CLI\CVSFiles\drive_VPD.csv
```

viewFibreChannel

> **Note:** All view commands generate a `.csv` by adding the **> *any_name_you_want*.csv** parameter at the end of the command.

Use the **viewFibreChannel** command to show the Fibre Channel settings for both ports in each drive. This command is useful to view the configuration of the ports for the Fibre Channel switch and to ensure that the ports are configured correctly to be compatible with the type of network in use.

The command uses the following syntax:

```
viewFibreChannel
```

The **viewFibreChannel** command does not require any parameters. When the **viewFibreChannel** command runs successfully, the system returns output to your display. The output file sorts the information by drive name, as shown in Example 5-103.

Example 5-103 The viewFibreChannel command

```
viewFibreChannel
       Drive,Location(F,C,R),Logical Library,  Type,  Port(1,2), Link
Status,Configured Link Speed,Configured Topology,Actual Link Speed,Actual Topology
0000013B008E, F1, C4, R1,Redbooks_custom,  3592-E07,  13,       No light,
Auto,  Auto (L Port),          Auto,        L Port
            ,           ,               ,          ,  77,       No light,
Auto,  Auto (L Port),          Auto,        L Port
0000013B006F, F1, C4, R4,Redbooks_custom,  3592-E07,  16,       No light,
Auto,  Auto (L Port),          Auto,        L Port
            ,           ,               ,          ,  80,       No light,
Auto,  Auto (L Port),          Auto,        L Port
```

The **viewFibreChannel** command can include a parameter that instructs the system to generate a .csv file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewFibreChannel > *any_name_you_want*.csv

The system generates the .csv file and stores it in the directory that contains the TS4500CLI.jar file. No message is returned to your display. To store the .csv file in another directory, specify the name of the directory in the command, as shown in Example 5-104.

Example 5-104 The viewFibreChannel command with a .csv file

```
viewFibreChannel > C:\CLI\CVSFiles\fibre_channel_settings.csv
```

viewIoStation

Use the **viewIoStation** command to view a summary of information for all cartridges in the I/O station. The command uses the following syntax:

```
viewIoStation
```

The **viewIoStation** command does not require any parameters. It lists the high-level information for all cartridges in the I/O station. When the **viewIoStation** command runs successfully, the system returns output to your display, as shown in Example 5-105.

Example 5-105 The viewIoStation command

```
viewIoStation
Volume Serial,Logical Library,Element Address,Media
Type,Location(F,C,R),Encryption
M000NTJL, Logical Library 1, 1046, JAG, I/O Slot(F1,C5,R4,T1), Not Encrypted
M001NTJL, Logical Library 1, 1294, JAG, I/O Slot(F1,C6,R29,T0), Not Encrypted
```

Use the **viewIoStation** command (and all **view** commands) to instruct the system to generate a .csv file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewIoStation > *any_name_you_want*.csv

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-106.

Example 5-106 The viewIoStation command with a .csv file

```
viewIoStation > C:\CLI\CVSFiles\viewIoStation.csv
```

viewISCSI

Use the **viewISCSI** command to configuration settings for a specific Ethernet drive.

The command uses the following syntax:

```
viewISCSI -f# -c# -r#
```

The number sign (#) variables specify the location of the drive, and have the following values:

-f# Frame number of the drive
-c# Column number of the drive
-r# Row number of the drive

When the **viewISCSI** command runs successfully, the system returns output to your display, as shown in Example 5-107.

Example 5-107 The viewiCSI command

```
viewISCSI -f4 -c3 -r1
iSCSI name:     naa.032165478954
iSCSI alias:    Drive05
DHCP:           Disabled
Port 0 address: 9.1.2.3
Port 1 address: 9.1.2.4
Gateway:        9.1.0.100
Subnet mask:    255.255.255.0
```

viewKeyLabelMapping

Use the **viewKeyLabelMapping** command to view the list of available key labels. The command uses the following syntax:

```
viewKeyLabelMapping
```

When the **viewKeyLabelMapping** command runs successfully, the system returns output to the display, as shown in Example 5-108.

Example 5-108 The viewKeyLabelMapping command

```
viewKeyLabelMapping
 Index,      Map From Key Label,     Key Mode,       Map To Key Label
    1,              REDB,          Wrapped-Hash,         END
    2,              REDB1,          Direct-Default-Set,  0
    3,              REDB1,         Direct-Specific,      Red1
```

The values in the Index column are referenced when you run the **editKeyLabelMapping** and **deleteKeyLabelMapping** commands.

viewLibraryVPD

Use the **viewLibraryVPD** command to show VPD for the library. This command is useful if an administrator or service engineer must see the model or serial number or other information specific to the library. This command requires a service user to run.

The command uses the following syntax:

```
viewLibraryVPD
```

When the **viewLibraryVPD** command runs successfully, the system returns output to your display, as shown in Example 5-109.

Example 5-109 The viewLibraryVPD command

```
viewLibraryVPD
       Location,       Machine Type,     Model,       Serial Number,Media Type
        Frame 1,               3584,       D25,             78CAFFF,      3592
        Frame 2,               3584,       L25,             78AAFFD,      3592
        Frame 3,               3584,       D55,             78DFFDS       LTO
```

The **viewLibraryVPD** command (and all **view** commands) can include a parameter that instructs the system to generate a .csv file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewLibraryVPD > *any_name_you_want*.csv

The system generates the .csv file and stores it in the directory that contains the TS4500CLI.jar file. No message is returned to your display. To store the .csv file in another directory, specify the name of the directory in the command, as shown in Example 5-110.

Example 5-110 The viewLibraryVPD command with a .csv file

```
viewLibraryVPD > C:\CLI\CVSFiles\loglib_VPD.csv
```

viewLogicalLibraries

Use the **viewLogicalLibraries** command to list the high-level information for all logical libraries. If you must list detailed information for a specific library, use the **viewLogicalLibraryDetails** command.

The command uses the following syntax:

```
viewLogicalLibraries
```

The **viewLogicalLibraries** command does not require any parameters. It lists the information for all logical libraries. When this command runs successfully, the system returns output to your display, as shown in Example 5-111.

Example 5-111 The viewLogicalLibraries command

```
viewLogicalLibraries
Name,    Type, Assigned Cartridges,  Virtual I/O cartridges,    Drives,Encryption
Method, Queued Exports,VOLSER Reporting (6/8 characters)
LTFS_redbooks, 3592, 23, 0, 4, No, Show, 8
Redbooks_custom, 3592, 235, 0, 2, Application managed, Show, 8
SLS_JAG1, 3592, 197, 0, 3, No, Show, 8
```

> **Note:** The Queued Exports information and the VOLSER Reporting information are available only through the **viewLogicalLibraryDetails** command at the CLI. This information is not available through the TS4500 management GUI.

The **viewLogicalLibraries** command can include a parameter that instructs the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewLogicalLibraries > *any_name_you_want*.csv

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-112.

Example 5-112 The viewLogicalLibraries command

```
viewLogicalLibraries > C:\CLI\CVSFiles\viewLogicalLibraries.csv
```

viewLogicalLibraryDetails

Use the **viewLogicalLibraryDetails** command to show the name, type, maximum number of cartridges, number of virtual I/O slot addresses, and number of drive addresses in a specified logical library.

The command uses the following syntax:

viewLogicalLibraryDetails "logical library name"

> **Important:** Parameters for the **viewLogicalLibraryDetails** command must be enclosed in double quotation marks.

When the **viewLogicalLibraryDetails** command runs successfully, the system returns output to your display, as shown in Example 5-113.

Example 5-113 The viewLogicalLibraryDetails command

```
viewLogicalLibraryDetails "library 1"
Name, Type, Max Cartridges, Virtual I/O slot addresses,Drive Addresses,
Library1, 3592, 1000, 130(768-897), 68(257-321)
```

The **viewLogicalLibraryDetails** command (and all **view** commands) can include a parameter that instructs the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewLogicalLibraryDetails > *any_name_you_want*.csv

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-114.

Example 5-114 The viewLogicalLibraryDetails command with a .csv file

```
viewLogicalLibraryDetails "library 1" > C:\CLI\CVSFiles\loglib_details.cs
```

viewMacAddress

Use the **viewMacAddress** command to look up the MAC addresses for a drive's Ethernet ports.

The command uses the following syntax:

```
viewMacAddress -f# -c# -r#
```

The number sign (#) variables specify the location of the drive, and have the following values:

-f# Frame number of the drive
-c# Column number of the drive
-r# Row number of the drive

When the **viewMacAddress** command runs successfully, the system returns output to your display, as shown in Example 5-115.

Example 5-115 The viewMacAddress command

```
viewMacAddress -f4 -c3 -r1
Port 0 MAC Address:  00:1a:64:eb:04:83
Port 1 MAC Address:  00:1a:64:eb:04:84
```

viewNodeCards

Use the **viewNodeCards** command to show a summary of information about the node cards.

The command uses the following syntax:

```
viewNodeCards
```

The **viewNodeCards** command does not require any parameters. It lists the high-level information for all node cards in the system. When the **viewNodeCards** command runs successfully, the system returns output to your display, as shown in Example 5-116.

Example 5-116 The viewNodeCards command

```
viewNodeCards
                 Card name,  Location,Part number,Serial number, Firmware version
                 XY Card A, Accessor A,  95P8618,   YN11MA42NF87,
                 XY Card B, Accessor B,  95P8618,   YN11MA57SGZD,
Accessor Controller Card A, Accessor A,  46X5961,   YN10MA42WFLT,
Accessor Controller Card B, Accessor B,  46X5961,   YN10MA5AH03L,
                     LCC A,    Frame 1,  12X5001,   Y1Y0103CK628, 1.8.0.1-C00
                     LCC A,    Frame 2,  12X5001,   Y1Y01055B01N, 1.8.0.1-C00
                     LCC A,    Frame 3,  12X5001,   Y1Y01077761G, 1.8.0.1-C00
```

Use the **viewNodeCards** command to instruct the system to generate a .csv file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewNodeCards > *any_name_you_want*.csv

The system generates the .csv file and stores it in the directory that contains the TS4500CLI.jar file. No message is returned to your display. To store the .csv file in another directory, specify the name of the directory in the command, as shown in Example 5-117.

Example 5-117 The viewNodeCards command with a .csv file

```
viewNodeCards > C:\CLI\CVSFiles\node_cards.csv
```

viewOfflineComponents

Use the **viewOfflineComponents** command to show a list of offline slots.

The command uses the following syntax:

```
viewOfflineComponents
```

The **viewOfflineComponents** command does not require any parameters. It lists all offline storage slots. When the **viewOfflineComponents** command runs successfully, the system returns output to your display, as shown in Example 5-118.

Example 5-118 The viewOfflineComponents command

```
viewOfflineComponents
Frame,     Column ,Row , Status
   1 ,      3     ,  1  , Offline
   1 ,      3     ,  2  , Offline
   1 ,      3     ,  3  , Offline
```

Use the **viewOfflineComponents** command to instruct the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewOfflineComponents > *any_name_you_want*.csv

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-119.

Example 5-119 The viewOfflineComponents command with a .csv file

```
viewOfflineComponents > C:\CLI\CVSFiles\Offline_Slots.csv
```

viewPasswordAndSessionPolicy

Use the **viewPasswordAndSessionPolicy** command to display the settings for GUI Sessions and the password characteristics set.

The command uses the following syntax:

```
viewPasswordAndSessionPolicy
```

When the **viewPasswordAndSessionPolicy** command runs successfully, the system returns output to the display, as shown in Example 5-120.

Example 5-120 The viewPasswordAndSessionPolity command

```
viewPasswordAndSessionPolicy

Automatic logout: 60
Password lock: 10 attempts
Automatic IMC (local GUI) login at power on: Disabled
Minimum number of characters: 8
Minimum number of uppercase characters: 1
Minimum number of lowercase characters: 1
Minimum number of numeric characters: 1
Minimum number of special characters: 0
Maximum identical consecutive characters: 2
Maximum password age: 90 day(s)
```

```
Minimum password age: 1 day(s)
Number of unique passwords before reusing: 8
```

viewRolePermissions

Use the **viewRolePermissions** command to display a list of all of the permissions for a specified role. The output shows a list of the web pages in the library that are available to that role, and the corresponding access level.

The command uses the following syntax:

```
viewRolePermissions role name filename.txt
```

The **filename.txt** variable is optional. When you specify a file name, you instruct the system to generate a .txt file that can be used with the **setRolePermissions** command. The system generates the .txt file and stores it in the directory that contains the TS4500CLI.jar file. The text file can be edited to change any of the access levels for the listed pages.

The following access levels are valid:

Read only The user has read-only access to the web page.
Modify The user can run the edit operation for the web page.
No Access The user cannot view the web page.

When the **viewRolePermissions** command runs successfully, the system returns output to your display, as shown in Example 5-121.

Example 5-121 The viewRolePermissions command

```
viewRolePermissions Redbooks name role.txt
                             Action,   Access Level
                                AIT,      No Access
                         Cartridges,         Modify
      Cartridges by Logical Library,         Modify
               Cleaning Cartridges,         Modify
                     Debug Messages,      No Access
                             Drives,         Modify
          Drives by Logical Library,         Modify
                Email Notifications,         Modify
                   Email Recipients,         Modify
                Encryption Internal,         Modify
             Encryption Key Manager,         Modify
                     Ethernet Ports,         Modify
                             Events,         Modify
                Library Information,         Modify
                 Licensed Functions,         Modify
                  Logical Libraries,         Modify
            Management GUI Behavior,         Modify
                     Master Console,      No Access
                     Password Rules,         Modify
                   Position Control,      No Access
              Remote Authentication,         Modify
                              Roles,         Modify
                  SNMP Destinations,         Modify
                 SNMP Notifications,         Modify
                         Scan Speed,      No Access
                Secure Socket Layer,         Modify
```

```
                  Service Port,      No Access
         Syslogs Notifications,         Modify
                        System,         Modify
          System Date and Time,         Modify
                         Tasks,         Modify
                         Users,         Modify
                 VOLSER Ranges,         Modify
  VOLSER Ranges by Logical Library,     Modify
```

viewRoles

Use the **viewRoles** command to display a list of all of the defined roles in the library.

The command uses the following syntax:

```
viewRoles
```

The **viewRoles** command does not require any parameters. It instructs the system to display information for all roles in the library. When the **viewRoles** command runs successfully, the system returns output to your display, as shown in Example 5-122.

Example 5-122 The viewRoles command

```
viewRoles
                Role,Mapped users
       Administrator,        5
           Superuser,        0
             Service,        2
             Monitor,        1
             Custom1,        0
            Redbooks,        0
           RemoteAcc,        0
```

The **viewRoles** command can include a parameter that instructs the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated.

The command uses the following syntax:

```
viewRoles > any_name_you_want.csv
```

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-123.

Example 5-123 The viewRoles command with a .csv file

```
viewRoles > C:\CLI\CVSFiles\Roles.csv
```

viewSnapshots

Use the **viewSnapshots** command to display a list of all available snapshots in the library. The files in the list can be used with the **downloadSnapshot** command.

The command uses the following syntax:

```
viewSnapshots
```

The **viewSnapshots** command does not require any parameters. It instructs the system to display information for all available snapshots in the library. When the **viewSnapshots** command runs successfully, the system returns output to your display, as shown in Example 5-124.

Example 5-124 The viewSnapshots command

```
viewSnapshots
1 - TS4500_SS_FA010_20140515124107.zip
2 - TS4500_SS_FA010_20140519105428.zip
3 - TS4500_SS_FA010_20140519105917.zip
4 - TS4500_SS_FA010_20140519110938.zip
5 - TS4500_SS_FA010_20140520213407.zip
6 - TS4500_SS_FA010_20140520213706.zip
```

The **viewSnapshots** command (and all **view** commands) can include a parameter that instructs the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated:

`viewSnapshots > any_name_you_want.csv`

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-125.

Example 5-125 The viewSnapshots command with a .csv file

```
viewSnapshots > C:\CLI\CVSFiles\Snapshots.csv
```

viewStorageCapacity

Use the **viewStorageCapacity** command to view the licensed media types. It provides more information about how each slot type (LTO / JAG) is licensed and how many slots are in use.

The command uses the following syntax:

`viewStorageCapacity`

The **viewStorageCapacity** command does not require any parameters. When the command runs successfully, the system returns output to your display, as shown in Example 5-126.

Example 5-126 The viewStorageCapacity command

```
viewStorageCapacity
[{"librarymediatype":"2","licensed_jag_used":"453","licensed_lto":"970","licensed_
jag":"1210","licensed_lto_used":"5"}]
```

viewSystemSummary

Use the **viewSystemSummary** command to view the physical library system summary. The **viewSystemSummary** command is useful for users who want to view information about each frame, or see a summary of the total storage capacity.

The command uses the following syntax:

`viewSystemSummary`

The **viewSystemSummary** command does not require any parameters. It lists the information about each frame in ascending order by frame number. When the **viewSystemSummary** command runs successfully, the system returns output to your display, as shown in Example 5-127.

Example 5-127 The viewSystemSummary command

```
viewSystemSummary
                    Frame:               1
                    State:     Door Closed
               Media Type:            3592
          Data Cartridges:               0
            Storage Slots:             660
                   Drives:               2
                      MTM:         3584L25
                      S/N:         13FA010
-------------------------------------------
                    Frame:               2
                Accessors:     OK - Online
                    State:     Door Closed
               Media Type:            3592
          Data Cartridges:              15
            Storage Slots:             875
                   Drives:               0
                      MTM:         3584D25
                      S/N:         13FA107
-------------------------------------------
      Total storage slots:            1535
  Total storage cartridges              15
           Total IO slots:               0
       Total IO cartridges              13
             Total Drives:               2
              Total Frames               2
```

The **viewSystemSummary** command can instruct the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

`viewSystemSummary > any_name_you_want.csv`

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-128.

Example 5-128 The viewSystemSummary command with a .csv file

```
viewSystemSummary > C:\CLI\CVSFiles\system_summary.csv
```

viewSystemSummaryDetails

Use the **viewSystemSummaryDetails** command to view detailed information about a specified frame. The **viewSystemSummaryDetails** command does not instruct the system to list summary information about all frames in the library. If you must view summary information about all frames in the library, use the **viewSystemSummary** command.

The command uses the following syntax:

`viewSystemSummaryDetails frame#`

The pound sign (#) variable in the command specifies the frame number of the frame for which you want more information. If the input frame is not available in the library, the system displays an error message.

When the **viewSystemSummaryDetails** command runs successfully, the system returns output to your display, as shown in Example 5-129.

Example 5-129 The viewSystemSummaryDetails command

```
viewSystemSummaryDetails 2
                    Frame:             2
                    State:     Door Closed
               Media Type:          3592
          Data Cartridges:           206
            Storage Slots:           550
                   Drives:             4
      Upper IO cartridges:             0
      Upper IO door slots:            16
 Upper IO door media type:          3592
      Lower IO cartridges:             0
      Lower IO door slots:            16
 Lower IO door media type:          3592
                      MTM:       3584D25
                      S/N:       78D3746
```

The **viewSystemSummaryDetails** command can instruct the system to generate a .csv file from the output. Enter a name, as a parameter to the command for the file that is generated, as shown in the following example:

viewSystemSummaryDetails > *any_name_you_want*.csv

The system generates the .csv file and stores it in the directory that contains the TS4500CLI.jar file. No message is returned to your display. To store the .csv file in another directory, specify the name of the directory in the command, as shown in Example 5-130.

Example 5-130 The viewSystemSummaryDetails command with a .csv file

```
viewSystemSummaryDetails > C:\CLI\CVSFiles\system_summary_Detail.csv
```

Note: The system does not display any information about accessors if the specified frame does not have accessors.

viewUsers

Use the **viewUsers** command to display a list of all of the users in the library.

The command uses the following syntax:

viewUsers

The **viewUsers** command does not require any parameters. It instructs the system to display information for all users in the library. When the **viewUsers** command runs successfully, the system returns output to your display, as shown in Example 5-131 on page 388.

Example 5-131 The viewUsers command

```
viewUsers
    Name,Locked,     State,           Role, Email,      Last login
   admin, , Disconnected,  Administrator,,Wednesday, May 14, 2014, 06:20:16 AM MST
  admin2, , Disconnected,  Administrator,, Tuesday, April 8, 2014, 07:02:48 AM MST
LocalGUI, , Disconnected,        Monitor,,  Tuesday, May 20, 2014, 02:43:50 PM MST
 Redbooks, , Connected,  Administrator,,Wednesday, May 21, 2014, 09:33:06 AM MST
 Service, , Disconnected,        Service,,  Tuesday, May 20, 2014, 02:42:36 PM MST
```

The **viewUsers** command can include a parameter that instructs the system to generate a .csv file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

viewUsers > *any_name_you_want*.csv

The system generates the .csv file and stores it in the directory that contains the TS4500CLI.jar file. No message is returned to your display. To store the .csv file in another directory, specify the name of the directory in the command, as shown in Example 5-132.

Example 5-132 The viewUsers command with a .csv file

```
viewUsers  > C:\CLI\CVSFiles\viewUsers.csv
```

viewUtilThreshold

Use the **viewUtilThreshold** command to see more information about the defined Capacity Utilization Threshold and the Dual Accessor Utilization Threshold.

If the number of cartridges exceeds the defined capacity utilization threshold, the Capacity POD in the monitor system page turns yellow. If this scenario occurs, remove the cartridges, increase the capacity threshold, or purchase more capacity.

The dual accessor utilization threshold (-daut) option is applicable to dual accessor libraries only with elastic capacity enabled. It determines the threshold in which cartridges are moved to accessor service areas in the end frames. The default is 98%.

The command uses the following syntax:

viewUtilThreshold

The viewUtilThreshold command does not require any parameters. When the viewUtilThreshold command runs successfully, the system returns output to the display, as shown in Example 5-133.

Example 5-133 The viewUtilThreshold command

```
viewUtilThreshold
Capacity Utilization Threshold: 99.0
Dual Accessor Utilization Threshold: 98
```

viewVolserRanges

Use the **viewVolserRanges** command to list all of the VOLSER ranges in the physical library. To list all of the VOLSER ranges that are assigned to each logical library, use the **viewVolserRangesByLL** command.

The command uses the following syntax:

```
viewVolserRanges
```

The **viewVolserRanges** command does not require any parameters. It lists the information for all of the VOLSER ranges in the physical library. When the **viewVolserRanges** command runs successfully, the system returns output to your display, as shown in Example 5-134.

Example 5-134 The viewVolserRanges command

```
viewVolserRanges
  Volser start - Volser end,Logical Library,  Media Type,Number of cartridges
            DF0000 - DF9999,Redbooks_custom,     3592,          0
            JK0000 - JK0009,    SLS_JAG1,        3592,          2
            RBA000 - RBA999,LTFS_redbooks,       3592,          0
            RBD000 - RBD099,Redbooks_custom,     3592,          0
```

The **viewVolserRanges** command can include a parameter that instructs the system to generate a `.csv` file from the output. Enter a name, as a parameter to the command, for the file that is generated, as shown in the following example:

```
viewVolserRanges > any_name_you_want.csv
```

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-135.

Example 5-135 The viewVolserRanges command with a .csv file

```
viewVolserRanges > C:\CLI\CVSFiles\volser_ranges.csv
```

viewVolserRangesByLL

Use the **viewVolserRangesByLL** command to list all of the VOLSER ranges for each logical library in the physical library. To list all of the VOLSER ranges for the physical library, use the **viewVolserRanges** command.

The command uses the following syntax:

```
viewVolserRangesByLL
```

The **viewVolserRangesByLL** command does not require any parameters. It lists the information for all of the VOLSER ranges in each logical library in the physical library. When the **viewVolserRangesByLL** command runs successfully, the system returns output to your display, as shown in Example 5-136.

Example 5-136 The viewVolserRangesByLL command

```
viewVolserRangesByLL
Logical Library,  Cartridges,  Media Type
  LTFS_redbooks,            0,          3592
RBA000 - RBA999,            0
Redbooks_custom,            0,          3592
DF0000 - DF9999,            0
RBD000 - RBD099,            0
       SLS_JAG1,            0,          3592
 K0000 - JK0009,            2
```

Use the **viewVolserRangesByLL** command to instruct the system to generate a `.csv` file from the output. Enter a name as a parameter to the command for the file that is generated, as shown in the following example:

`viewVolserRangesByLL >` *`any_name_you_want`*`.csv`

The system generates the `.csv` file and stores it in the directory that contains the `TS4500CLI.jar` file. No message is returned to your display. To store the `.csv` file in another directory, specify the name of the directory in the command, as shown in Example 5-137.

Example 5-137 The viewVolserRangesByLL command with a .csv file

```
viewVolserRangesByLL > C:\CLI\CVSFiles\volser_ranges_by_ll.csv
```

6

TS4500 REST API

This chapter describes the TS4500 REST API.

The TS4500 REST API can be used as an alternative to the TS4500 CLI as a method of obtaining information from and sending commands to the library.

As of this writing, only the in-band method for sending REST API commands is available. In-band means that all commands are sent to the library by way of a Control Path drive from an attached host. Because this process uses SCSI commands, it is called *REST over SCSI* (RoS).

RoS allows library administrators to obtain status information about library components by sending REST API (by using SCSI Write Buffer commands) and receiving HTTP responses (by using Read Buffer commands). If RoS is enabled (the default setting), the library accepts and responds to commands. If disabled, commands are rejected. For more information about checking the setting, see 4.3, “Settings” on page 193.

The commands can be sent to the library by using the `sg3_utils` package (Linux) or ITDT.

This chapter includes the following topics:

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

6.1 REST API Overview

This section describes REST API commands including definitions, formats, endpoints, parameters, and examples. We describe how to use the in-band method for sending REST API commands and receiving HTTP responses by using SCSI Write Buffer and Read Buffer commands. Also, we show how to use IBM Tape Diagnostic Tool (ITDT) for RoS.

6.1.1 Commands

The REST API commands can be one of the following:

- `GET` commands that request information.
- `POST` commands that perform a longer task (such as an inventory) or an action (such as cleaning a drive).
- `PATCH` commands that perform an immediate action (such as setting the time).

These commands are issued along with an endpoint (the resource) and, if required, parameters.

6.1.2 Endpoints

The following endpoints are available (with library firmware version 1.7.0.0 and later):

- Accessors:
  - `GET /v1/accessors`
  - `GET /v1/accessors/<location>`
- Cleaning Cartridges:
  - `GET /v1/cleaningCartridges`
  - `GET /v1/cleaningCartridges/<volser>`
  - `POST /v1/cleaningCartridges/<internalAddress>/moveToIOStation`
- Data Cartridges
  - `GET /v1/dataCartridges`
  - `GET /v1/dataCartridges/<volser>`
  - `POST /v1/dataCartridges/<internalAddress>/moveToIOStation`
- Diagnostic Cartridges:
  - `GET /v1/diagnosticCartridges`
  - `GET /v1/diagnosticCartridges/<volser>`
  - `POST /v1/diagnosticCartridges/<volser>/moveToIOStation`
- Drives:
  - `GET /v1/drives`
  - `GET /v1/drives/<location>`
  - `GET /v1/drives/<sn>`
  - `POST /v1/drives/<location>/clean`
  - `PATCH /v1/drives/<location> {“use”: <“access” | “controlPath” | “verification”>}`
  - `POST /v1/drives/<location>/reset {“mode”: <“normal” | “hard”>}`
  - `PATCH /v1/drives/<location> {“beacon”: <“enabled” | “disabled”>}`
- Ethernet Ports:
  - `GET /v1/ethernetPorts`
  - `GET /v1/ethernetPorts/<location>`

- Events:
  - GET /v1/events
  - GET /v1/events/<ID>
  - GET /v1/events?<“after” | “before”>=<time>
  - GET /v1/events/<location>
- Frames:
  - GET /v1/frames
  - GET /v1/frames/<location>
- Library:
  - GET /v1/library
  - POST /v1/library/reset
  - GET /v1/library/saveConfig
  - PATCH /v1/library {“time”:<“time”>}
  - GET /v1/powerSupplies
  - GET /v1/slots
- Logs
  - GET /v1/logs
  - GET /v1/logs/<filename>
  - POST /v1/logs
  - GET /v1/logs/<filename>/export
- Node Cards:
  - GET /v1/nodeCards
  - GET /v1/nodeCards/<ID>
  - POST /v1/nodeCards/<ID>/reset
- Reports:
  - GET /v1/reports/drives
  - GET /v1/reports/drives/<location>
  - GET /v1/reports/drives/<sn>
  - GET /v1/reports/library
  - GET /v1/reports/accessors
  - GET /v1/reports/accessors/<location>
- Tasks:
  - GET /v1/tasks
  - GET /v1/tasks/<ID>
  - POST /v1/tasks {“type”: “inventoryTier0and1”, “location”: <“library” | “frame_F<f>”>}
  - POST /v1/tasks {“type”: “inventoryAllTiers”, “location”: <“library” | “frame_F<f>”>}
  - POST /v1/tasks {“type”: “calibrateLibrary”, “location”: “library”}
  - POST /v1/tasks {“type”: “calibrateFrame”, “location”: “frame_F<f>”}
  - POST /v1/tasks {“type”: “calibrateAccessor”, “location”: “accessor_A<“a”|”b”>”}
  - POST /v1/tasks {“type”: “testDrive”, “location”: “drive_F<f>C<c>R<r>”}
  - POST /v1/tasks {“type”: “startDriveService”, “location”: “drive_F<f>C<c>R<r>”}
  - POST /v1/tasks {“type”: “completeDriveService”, “location”: “drive_F<f>C<c>R<r>”}
  - POST /v1/tasks {“type”: “startAccessorService”, “location”: “accessor_A<“a”|”b”>”}
  - POST /v1/tasks {“type”: "completeAccessorService", “location”: accessor_A<“a”|”b”>”

Introduction of Work Items

Moving a cartridge was attempted in the past by using the “moveToIOSlot” POST actions that were added to the various cartridge type resources. However, that method did not consider the asynchronous nature of the move and treated it as a synchronous action.

A new resource called a workItems is a new resource that manages the movement of cartridges. These work items make up the work queue that is run by the library.

The intention of workItems is to provide a more-capable (and transparent) API that can replace the use of SCSI Move Medium. To that end, when SCSI and web move commands differ in behavior, the intention is that the **REST** command follows the SCSI convention; for example, maintaining the logical library assignment on a move to an I/O Station.

As of this writing, the work queue cannot be queried, although this function is being considered in a future product (long-term roadmap) to allow users to track the completion of this movement and see a history of movement by the user.

As a mid-term solution, a new lastMoveResult attribute is added for all cartridge types.

For now, the user queries the cartridge that is being moved and poll on it to wait for the location attribute to change indicating it has reached its destination.

The following workitems commands and the deprecated commands are available:

- Work Items:
  - POST /v1/workItems [{"type": "moveToDrive", "cartridge": <volser>,"sourceInternalAddress": <internalAddress>, "destinationLocation": <location>, "destinationSN": < serialNumber >}]
  - POST /v1/workItems [{"type": "moveToSlot", "cartridge": <volser>,"sourceInternalAddress": <internalAddress>}]
  - POST /v1/workItems [{"type": "moveToIOStation", "cartridge": <volser>, "sourceInternalAddress": <internalAddress>}]
- Deprecated URL endpoints:
  - POST /v1/cleaningCartridges/<internalAddress>/moveToIOStation
  - POST /v1/dataCartridges/<internalAddress>/moveToIOStation
  - POST /v1/diagnosticCartridges/<internalAddress>/moveToIOStation

The TS4500 RoS commands and the TS4500 CLI equivalents are listed in Table 6-1.

Table 6-1 TS4500 RoS commands versus TS4500 CLI commands

REST over SCSI commands	CLI commands
GET /v1/library	viewSystemSummary
GET /v1/accessors	viewAccessor
GET /v1/accessors/<location>	viewAccessor
GET /v1/cleaningCartridges	viewCleaningCartridges
GET /v1/cleaningCartridges/<volser>	viewCleaningCartridges
GET /v1/dataCartridges	viewDataCartridges
GET /v1/dataCartridges/<volser	N/A
GET /v1/diagnosticCartridges	N/A
GET /v1/diagnosticCartridges/<volser>	N/A
POST /v1/diagnosticCartridges/<volser>/moveToIOStation	N/A
GET /v1/drives	viewDriveSummary

REST over SCSI commands	CLI commands
GET /v1/drives/<location>	viewDriveDetails
GET /v1/drives/<sn>	viewDriveDetails
POST /v1/drives/<location>/clean	cleanDrive
PATCH /v1/drives/<location> {"use": <"access" | "controlPath" | "verification">}	setDriveUse
POST /v1/drives/<location>/reset {"mode": <"normal" | "hard">}	resetDrive
PATCH /v1/drives/<location> {"beacon": <"enabled" | "disabled">}	N/A
GET /v1/ethernetPorts	N/A
GET /v1/ethernetPorts/<location>	N/A
GET /v1/events	downloadEvents
GET /v1/events/<ID>	downloadEvents
GET /v1/frames	viewSystemSummary
GET /v1/frames/<location>	viewSystemSummaryDetails
GET /v1/library	N/A
POST /v1/library/reset	resetNodeCards
GET /v1/library/saveConfig	saveConfiguration
PATCH /v1/library {"time":<time>}	setLibraryTime
GET /v1/logs	N/A
GET /v1/logs/<filename>	N/A
POST /v1/logs	downloadLog
GET /v1/logs/<filename>/export	downloadLog
GET /v1/nodeCards	viewNodeCards
GET /v1/nodeCards/<ID>	viewNodeCards
POST /v1/nodeCards/<ID>/reset	resetNodeCards
GET /v1/powerSupplies	N/A
GET/v1/reports/drives	N/A
GET /v1/reports/drives/<location>	N/A
GET /v1/reports/drives/<sn>	N/A
GET/v1/reports/library	N/A
GET/v1/reports/accessors	N/A
GET/v1/reports/accessors/<location>	N/A
GET /v1/slots	N/A
GET /v1/tasks	N/A
GET /v1/tasks/<ID>	N/A

REST over SCSI commands	CLI commands
POST /v1/tasks {“type”: ”inventoryTier0and1”, “location”: <”library” | ”frame_F<f>”>}	startInventory
POST /v1/tasks {“type”: ”inventoryAllTiers”, “location”: <”library” | ”frame_F<f>”>}	startInventory
POST /v1/tasks {“type”: ”startDriveService”, “location”: ”drive_F<f>C<c>R<r>”}	startDriveService
POST /v1/tasks {“type”: ”completeDriveService”, “location”: ”drive_F<f>C<c>R<r>”}	completeDriveService
POST /v1/tasks {“type”: ”calibrateAccessor”, “location”: ”accessor_A<”a”|”b”>”}	startCalibration
POST /v1/tasks {“type”: ”startAccessorService”, “location”: ”accessor_A<”a”|”b”>”}	N/A
POST /v1/tasks {“type”: ”completeAccessorService”, “location”: ”accessor_A<”a”|”b”>”}	N/A
POST /v1/tasks {“type”: “testDrive”, “location”: ”drive_F<f>C<c>R<r>”}	N/A
POST /v1/workItems [{"type": "moveToSlot", "cartridge": <volser>,"sourceInternalAddress": <internalAddress>}]	N/A
POST /v1/workItems [{"type": "moveToIOStation", "cartridge": <volser>, "sourceInternalAddress": <internalAddress>}]	N/A
POST /v1/workItems [{"type": "moveToDrive", "cartridge": <volser>,"sourceInternalAddress": <internalAddress>, "destinationLocation": <location>, "destinationSN": < serialNumber >}]	N/A

6.1.3 Parameters

Some commands include optional parameters; for example, <location> in the **GET /v1/drives/<location>** command. The parameters can be classed as Path Parameters, Request Body Parameters, or Query Parameters. The following parameters are available:

- Accessors

  The location string describes the unique identifier for the accessor with a value of **“**accessor_Aa**”** or **“**accessor_Ab**”**.

- Cleaning cartridges

  The volser string represents the volume serial number or bar code that uniquely identifies the cartridge to a host; for example, CLN137JA. If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

- Data Cartridges

  The volser string represents the volume serial number or bar code that uniquely identifies the cartridge to a host; for example, SGP293JB. If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

- Diagnostic Cartridges

  The volser string represents the volume serial number or bar code that uniquely identifies the diagnostic cartridge; for example, CE 882JA. If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

- Drives:
  - location: A string that represents the unique location of the tape drive. This identifier is also the unique identifier for the drive. The format “drive_F<f>C<c>R<r>” is used in which "f" is the frame, "c" is the column, and "r" is the row.
  - sn: A string that represents the serial number of the drive. This string also can be used as a unique identifier for the tape drive. The format “0000078XXXXX” is used in which *78XXXXX* is the 7-digit serial number of the drive.
  - use: A string that represents the assigned use this tape drive was given within the logical library. The following values are included:
    - access: Data access drive
    - controlPath: Data access and control path drive
    - verification: Media verification drive
  - mode: An optional string that represents the type of reset to perform on the drive. If this string is not provided, normal is used. Possible values include:
    - normal: A command is sent to the drive and indicates it should reboot itself
    - hard: The library removes power to the drive temporarily to force a reset
  - beacon: A string that indicates the current state of the drive beacon LED. This beacon can be seen from the rear of the drive. Values include:
    - enabled (LED is flashing)
    - disabled (LED is off).

- Ethernet Ports

  The location string represents the unique location of this port. The format of this field is “ethernetPort_F<f>P<p>” with possible values for <p> of "a", "b", "imc", "tssc", or "service".

- Events:
  - ID: A decimal integer that is the unique identifier of the event in the library.
  - after (optional): A string that represents the time after which events are viewed. Any event with a time before this time is not returned. Time format is in ISO 8601 format of “YYYY-MM-DDThh:mm:ss+-hhmm” or “YYYY-MM-DDThh:mm:ss”; for example, “2018-09-17T23:02:00”. If the time zone is omitted, the current system’s current time zone is used.
  - before (optional): A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of “YYYY-MM-DDThh:mm:ss+-hhmm” or “YYYY-MM-DDThh:mm:ss”; for example, "2018-09-17T23:02:00". If the time zone is omitted, the system’s current time zone is used.
  - location (optional): A string that represents the hardware component in which the event originated. Hardware components include:
    - accessor_A<"a"|"b"> (for example, accessor_Ab)
    - gripper_A<"a"|"b">G<"1"|"2"> (for example, gripper_AbG1)
    - column_F<f>C<c> (for example, column_F1C3)
    - frame_F<f> (for example, frame_F1)
    - fiducial_F<f>C<c>L<"t"|"b"> (for example, fiducial_F1C3IOuLt)
    - fiducial_F<f>IO<"u"|"l">L<"t"|"b"> (for example, fiducial_F1IOuLb)

    - drive_F<f>C<c>R<r> (for example, drive_F1C3R23)
    - slot_ F<f>C<c>R<r>T<t> (for example, slot_F1C3R23T0)
    - ioStation_F<f>IO<"u"|"l"> (for example, ioStation_F1IOu)
    - ethernetPort_F<f>P<p> (for example, ethernetPort_F2Pa)
    - frameSide_F<f>S<"A"|"B"> (for example, frameSide_F1Sa)
    - position_F<f>C<c>R<r> (for example, position_F1C5R7)

- Frames

  The location string represents the unique location of the frame in the library. It also reflects the type of frame. The format is “frame_F<f>” where ‘f’ is the frame position.

- Library

  The time parameter is the current date and time that are set on the library. All dates and times that are returned by the library are in relation to this time. Time format is in ISO 8601 format of **“**YYYY-MM-DDThh:mm:ss+-hhmm**”**; for example, **“**2018-09-17T23:02:00-0700**”**.

- Logs

  The filename parameter is the name of the file that is created or while being created. This file name contains the name of the library, the date and time it was taken, and other important service information.

- Node Cards

  The ID parameter is the decimal integer that is the unique identifier of the node card in the library.

- Reports

  The reports contain usage history and other data for resources in the library.

  Using a time stamp reports can be requested after or before a specific time. The time format is in ISO 8601 "YYYY-MM-DDThh:mm:ss+-hhmm", as shown in the following example,

  2020-03-12T17:02:00-0700

  The following query parameters are available:

  – Drives

    - after (optional): A string that represents the time after which events are viewed. Any event with a time before this time is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-12T17:02:00". If the time zone is omitted, the current system’s current time zone is used.

    - before (optional): A string that represents the time before which events should be viewed. Any event with a time after this time is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-12T13:02:00". If the time zone is omitted, the current systems current time zone is used.

  – Library

    The library report contains usage history of the library including mounts, inserts, ejects, moves, host I/O, and environmental data. These reports are taken on intervals of 1 hour and kept for 1 year. However, you can offload only 1 week of data at a time per GET command to ensure that the query response size does not get too large. If you do not include query parameters to limit the time frame, all reports for the last week are returned.

  The following (optional) command parameters are available:

  - after: A string that represents the time after which events are viewed. Any event with a time before this time is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-12T17:02:00". If the time zone is omitted, the current system’s current time zone is used.
  - before: A string that represents the time before which events should be viewed. Any event with a time after this is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-12T13:02:00". If the time zone is omitted, the current systems current time zone is used.

- Accessors

  The following (optional) parameters are available:

  - after: A string that represents the time after which events are viewed. Any event with a time before this time is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-12T17:02:00". If the time zone is omitted, the current system’s current time zone is used.
  - before: A string that represents the time before which events should be viewed. Any event with a time after this is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-12T13:02:00". If the time zone is omitted, the current systems current time zone is used.

► Tasks:

- ID: A decimal integer that is the unique identifier of the task in the library.
- type: A string that represents the task type to run. Values can include:
  - inventoryTier0and1
  - inventoryAllTiers
  - startDriveService
  - completeDriveService
  - calibrateAccessor
  - startAccessorService
  - completeAccessorService
  - testDrive
- location: A string that represents the hardware component what the task is affecting. This string can be null if no hardware was associated with this task. Hardware components include:
  - library or frame_F<f> where “f” is the frame number of a single frame
  - drive_F<f>C<c>R<r> where “f” is the frame, “c” is the column, and “r” is the row
  - accessor_A<"a"|"b"> where “a” is accessor A and “b” is accessor B

6.2 Using sg3_utils (Linux)

The `sg3_utils` is part of the Linux package. The `sg3_utils` package includes utilities that send SCSI commands to devices.

Next, we describe the following in-band method for sending REST API commands (by using SCSI Write Buffer commands) and receiving HTTP responses by using Read Buffer commands:

1. We write the buffer with the information that we want to know. In this case, we want to ask for the library information to be written into the buffer.
2. After the buffer is written, we can read the data from the buffer. Whenever the data is read, the buffer is empty and must be rewritten.

This task is done by using the **`sg_write_buffer`** and **`sg_read_buffer`** commands being run in turn. This process can be done manually, by scripting, or other programming.

The following parameters also are required:

- Buffer ID for each REST API command is X'10'
- Buffer ID for the HTTP response to the most recent REST API command is X'11'
- Buffer Mode is 2 for all RoS commands
- Medium changer device name

The following parameters that are used in building the **`/usr/bin/sg_write_buffer`** command are listed in Table 6-2:

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg14<<<'GET /v1/library'
```

Table 6-2 /usr/bin/sg_write_buffer command

Parameters	Description
--id=0x10	Write buffer always has the id=0x10
--mode=2	Mode is always = 2
--raw	Request raw data
/dev/sg14	The medium changer device
<<<'GET /v1/library'	Request library endpoint information

After the buffer is written, we can read the data from the buffer. RoS buffers are volatile and after they are read once, they are empty.

The parameters that used in building the **`/usr/bin/sg_read_buffer`** command are listed in Table 6-3:

```
/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg14
```

Table 6-3 /usr/bin/sg_read_buffer command

Parameters	Description
--id=0x11	Read buffer always has the id=0x11
--mode=2	Mode is always = 2
--length=0x100000	Each read buffer response has a maximum transfer length of 1 MiB

Parameters	Description
/dev/sg14	The medium changer device

Each read buffer response includes a maximum transfer length of 1 MiB. To determine the remaining length of an HTTP response that was truncated to 1 MiB, use the value that is returned for the `Content-Length:` field in the header of the HTTP response (at Buffer Offset of 0). RoS buffers are volatile and are maintained separately for each control path in the library.

The SCSI Write Buffer (GET query) uses the following format:

```
sg_write_buffer --id=0x10 --mode=2 --raw /dev/$device<<<"GET /v1/library"
```

The SCSI Read Buffer (HTTP response to GET query) uses the following format:

```
sg_read_buffer --id=0x11 --mode=2 --length=$length --offset=0 --raw /dev/$device
```

The following query flows for large data transfers are available (greater than 1 MiB):

- `SCSI Write Buffer (GET query)`
- `SCSI Read Buffer (first portion of the HTTP response to GET query)`
- `--offset=0`
- `SCSI Read Buffer (next portion of the HTTP response to GET query)`
- `--offset=1`
- `..... (offset continues to increase until the end of the entire content)`
- `SCSI Read Buffer (final portion of the HTTP response to GET query)`

6.2.1 Example

Example 6-1 shows the command sequence to get the Library information. The output is unformatted.

Example 6-1 GET /v1/library command

```
[root@localhost scripts]# sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg73<<<"GET /v1/library"

[root@localhost scripts]# sg_read_buffer --id=0x11 --mode=2 --length=783  --raw
/dev/sg73
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 715

[{"name": "9_Framer", "status": "online", "totalCapacity": "9778",
"licensedCapacity": "5040", "totalCartridges": "2177", "assignedCartridges":
"557", "firmware": "1.6.0.1-B00", "sn": "78BB640", "time":
"2019-11-01T12:57:09-0600", "location": "Spectrum Lab", "address": "Campino 2200",
"city": "El Salto", "state": "Jalio", "country": "Mexico", "contact": "Mad Max",
"telephone": "0000", "secondaryTelephone": "4711", "secureCommunications":
"disabled", "autoEjectCleaningCartridges": "enabled", "elasticCapacity":
"enabled", "activeAccessors": "Accessor A only", "vioStatus": "enabled",
"nmaDetection": "disabled", "capacityUtilThresh": "99.0",
"dualAccessorUtilThresh": "98.0"}]
```

6.2.2 Endpoints

This section describes the endpoints and attributes that are defined in the REST API.

Accessors

Retrieve a list of accessors or a single accessor resource:

```
GET /v1/accessors
GET /v1/accessors/<location>
```

Path parameter

The `location` path parameter is a string that describes the unique identifier for the accessor with a value of `accessor_Aa` or `accessor_Ab`.

When the **`Get /v1/accessors/accessor_Ab`** command runs successfully, the system returns the output that is shown in Example 6-2 on page 402.

Example 6-2 Get /v1/accessor/accessor_Ab command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg14<<<'GET
/v1/accessors/accessor_Ab'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg14

[{"location": "accessor_Ab", "state": onlineStandby, "pivots": "11810",
"barCodeScans": "89980", "travelX": "3809", "travelY": "146518", "getsGripper1":
"8742", "putsGripper1": "8739", "getsGripper2": "1139", "putsGripper2": "1134"}]
```

The returned attribute values show the number of lifetime moves or meters traveled for each component. The “state” field returns the current state of the accessor. The possible states include the following values, which are listed in priority order:

- `inServiceMode`
- `noMovementAllowed`
- `noMotorPower`
- `bothGrippersFailed`
- `gripper1Failed`
- `gripper2Failed`
- `calibrating`
- `onlineStandby`
- `onlineActive`

Cleaning cartridges

Cleaning cartridges are used periodically to clean a drive. They have a limited number of uses.

The **`GET cleaning cartridges`** retrieves a list of system-managed cleaning cartridges or a single cleaning cartridge resource, as shown in the following examples:

```
GET /v1/cleaningCartridges
GET /v1/cleaningCartridges/<volser>
```

Path parameter

A `volser` is a string that represents the volume serial number or bar code that uniquely identifies the cartridge; for example, `CLN236JA`.

If duplicate VOLSERs are in the library, a query of this VOLSER return more than one cartridge.

When the **Get /v1/cleaningCartridges** command runs successfully, the system returns the output that is shown in Example 6-3.

Example 6-3 Get /v1/cleaningCartridges command

```
//usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg14<<<'GET 
/v1/cleaningCartridges'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg14

[{"volser": "CLN236JA", "state": "Normal", "cleansRemaining": "50", "location": 
"slot_F2C6R12T0", "mediaType": "3592", "mostRecentUsage": 
"2019-06-18T14:21:38-0600"}]
```

The following command attributes are available:

- `volser`: A string that represents the volume serial number or bar code that uniquely identifies the cartridge.
- `state`: A string that represents the current state of the cartridge. The following values are available (listed in priority order):
  - `failedVerification`
  - `assignmentRequired`
  - `uncertainBarcode`
  - `exportQueued`
  - `importing`
  - `verifying`
  - `normal`
- `cleansRemaining`: A string showing the number of cleans remaining.
- `location`: A string that describes the current location of the cartridge. The gripper location is shown only as part of a failure scenario if the cartridge is stuck in the gripper or as part of a tier 2 mount. Otherwise, a cartridge that is in transit from one location to another maintains its source location until after the move is completed.

  The format depends on the location of the cartridge, as shown in the following examples:

  - gripper_A<"a"|"b">G<"1"|"2"> (for example, `gripper_AbG1`)
  - drive_F<f>C<c>R<r> (for example, `drive_F1C3R23`)
  - slot_ F<f>C<c>R<r>T<t> (for example, `slot_F1C3R23T0`)
  - ioStationSlot_F<f>IO<"u"|"l">S<s> (for example, `ioStationSlot_F1IOuS4`)
- `mediaType`: A string that represents the media type that is supported by the tape drive. Supported values include `"3592"` and `"LTO"`.
- `mostRecentUsage`: A string that represents the last date and time this cartridge was mounted into a drive, or null if this is unknown or the cartridge is not yet mounted. Time format is in ISO 8601 format of `"YYYY-MM-DDThh:mm:ss+-hhmm"`; for example, `"2018-09-17T23:02:00-0700"`.

> Note: Any cartridge with an `Unknown` or null VOLSER is assumed by the library to be a data cartridge; therefore, they never appear in this query.

Data cartridges

Retrieves information about all host-accessible or unassigned data cartridges in the tape library. If a data cartridge is removed from the library or missing during an inventory scan, it will not appear in this list.

```
GET /v1/dataCartridges
GET /v1/dataCartridges/<volser>
```

Path parameter

A `string that represents` the volume serial number or bar code that uniquely identifies the cartridge to a host. For example, "`EZ1122L3`". If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

When the **`GET /v1/dataCartridges/volser`** command runs successfully, the system returns the output that is shown in Example 6-4 on page 404.

Example 6-4 GET /v1/dataCartridges command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg15<<<'GET
/v1/dataCartridges/VU0082JA'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg15
[{"volser": "VU0082JA", "state": "normal","accessible":"normal","location":
"slot_F1C8R4T0", "mediaType": "3592", "encrypted": null,"locked":"unlocked",
"mostRecentVerification": null, "mostRecentUsage": "2020-03-10T15:44:37-0600",
"logicalLibrary": null, "elementAddress": null}]
```

> **Note:** Any cartridge with an Unknown or null VOLSER is assumed by the library to be a data cartridge. Therefore, they never appear in this query.

The following command attributes are available:

- `volser`: The volser string represents the volume serial number or bar code that uniquely identifies the cartridge to a host; for example, EZ0137JA. If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.
- `state`: A string that represents the current state of the cartridge. Values include the following in priority order:
  - failedVerification
  - assignmentRequired
  - uncertainBarcode
  - exportQueued
  - importing
  - verifying
  - normal
- `location`: A string that describes the location of the cartridge.

> **Note:** The gripper location is shown only as part of a failure scenario if the cartridge is stuck in the gripper or as part of a tier 2 mount. Otherwise, a cartridge that is in transit from one location to another maintains its source location until after the move completes. The format depends on the location of the cartridge, as follows:
>
> - gripper_A<“a”|”b”>G<“1”|”2”> (for example, “gripper_AbG1”)
> - drive_F<f>C<c>R<r> (for example, “drive_F1C3R23”)
> - slot_ F<f>C<c>R<r>T<t> (for example, “slot_F1C3R23T0”)
> - ioStationSlot_F<f>IO<“u”|”l”>S<s> (for example, “ioStationSlot_F1IOuS4”)

- mediatype: A string that represents the media type that is supported by the tape drive. Supported values include “3592” and “LTO”.
- encrypted: A string that represents the state of encryption on this cartridge. If “yes”, the data on the cartridges was encrypted by the tape drive. If “no”, the data on the cartridge was not encrypted by the tape drive. If null, the cartridge is not yet mounted.
- mostRecentVerification: A string that represents the last date and time this cartridge was verified by using media verification in a drive that is assigned as a media verification drive. If media verification was not run on this cartridge, a null is returned.
- mostRecentUsage: A string that represents the last date and time this cartridge was mounted into a drive, or null if this is unknown or the cartridge was not yet mounted. Time format is in ISO 8601 format of “YYYY-MM-DDThh:mm:ss+-hhmm”. For example, “2020-03-17T13:02:00-0700”.
- logicalLibrary: A string that represents the name of the logical library to which this cartridge is assigned. If it is not assigned to a logical library, this value is null. If the cartridge does not belong to a bar code range, it is available to all logical libraries. In this case, the state reported is “importing” and this value returns null.
- elementAddress: An integer representing the element address of the cartridge drive. This value is reported to hosts by way of the Read Element Status SCSI command. This value is null if not yet reported to the host.

Diagnostic cartridges

Diagnostic cartridges are used to test tape drives, as shown in the following example:

```
GET /v1/diagnosticCartridges
GET /v1/diagnosticCartridges/<volser>
```

Path parameter

A volser string represents the volume serial number or bar code that uniquely identifies the diagnostic cartridge; for example, “DG 031L5”. If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

When the **Get /v1/diagnosticCartridges** command runs successfully, the system returns the output that is shown in Example 6-5.

Example 6-5 Get /v1/diagnosticCartridges command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg14<<<'GET
/v1/diagnosticCartridges'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg14
[{"volser": "DG 031L5", "state": "Normal", "location": "slot_F1C6R4T0",
"mediaType": "LTO", "mostRecentUsage": "2019-10-02T11:27:58-0600"}]
```

The attributes are the same as those attributes for Cleaning cartridges.

> **Note:** Any cartridge with an `Unknown` or null VOLSER is assumed by the library to be a data cartridge; therefore, they never appear in this query.

Moving diagnostic cartridges to I/O station

This command moves the specified diagnostic cartridge to the I/O Station. The export move is queued. Completion is indicated by the location attribute for the specified cartridge changing to an I/O station slot, as shown in the following example:

```
POST /v1/diagnosticCartridges/<volser>/moveToIOStation
```

Path parameter

A `volser` string represents the volume serial number or bar code that uniquely identifies the diagnostic cartridge; for example: `CE882JA`.

When the **`POST /v1/diagnosticCartridges/<volser>/moveToIOStation`** command runs successfully, the system returns the output that is shown in Example 6-6 on page 406.

Example 6-6 POST /v1/diagnosticCartridges/CE882JA/ moveToIOStation command

```
[root@TRON scripts]#/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'POST /v1/diagnosticCartridges/CE882JA/moveToIOStation'
```

The command prompt returns without comments. However, if we check the status of Diagnostic Cartridge CE882JA (**`GET /v1/diagnosticCartridges`**) before and after the command is run, it is moved to the VIO slot and later to the physical I/O station, as shown in the following examples:

- Before command execution (output truncated for better visibility):

  `“mediaType”: “3592”, “state”: “Normal”, “volser”: “CE 882JA”, “location”: “slot_F2C5R26T3”`

- After command execution (output truncated for better visibility):

  `“mediaType”: “3592”, “state”: “Virtual I/O”, “volser”: “CE 882JA”, “location”: “slot_F2C5R26T1”`

- After a short period:

  `“mediaType”: “3592”,”state”: “Normal”,”volser”: “CE 882JA”,”location”: “ioStationSlot_F2OuS1”`

Get Drives

Retrieve a list of drives or a single drive resource, as shown in the following example:

```
GET /v1/drives
GET /v1/drives/<location>
GET /v1/drives/<sn>
```

Path parameters

The following path parameters are available:

- `location`: A string that represents the unique location of the drive. This value also is the unique identifier for the drive. The format is “`drive_F<f>C<c>R<r>`” where “`f`” is the frame, “`c`” is the column, and “`r`” is the row in which the drive is installed.
- `sn`: A string that represents the serial number of the drive. This value can also be used as a unique identifier for the tape drive.

When the **Get /v1/drives** command runs successfully, the system returns output to your display, as shown in Example 6-7.

Example 6-7 Get /v1/drives command

```
[/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg70<<<'GET
/v1/drives/drive_F2C4R4'
/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70
[{"location": "drive_F2C4R4", "sn": "0000078PG204", "mediaType": "3592", "state":
"online", "operation": "empty", "mtm": "3592-60F", "bar code":
"11S02CE664YF1013000638", "interface": "fibreChannel", "interfaceMode":
"multi-mode", "logicalLibrary": "JAG6_2", "use": "controlPath", "firmware":
"559B", "encryption": "disabled", "wwnn": "500507604419ff0f", "elementAddress":
"264", "beacon": "disabled", "volser":
null}]
```

The following command attributes are available:

- `mediaType`: A string that represents the media type that is supported by the tape drive. Supported values include "`3592`" and "`LTO`".
- `state`: A string that represents the current state of the tape drive. The following values are available (in priority order):
  - `inServiceMode`
  - `restarting`
  - `initializing`
  - `unreachable`
  - `resetRequired`
  - `updating`
  - `cleaning`
  - `online`
- `operation`: A string that represents the operation that the tape drive is performing. If the tape drive is not in the `online` state, this value is null. The following values are available (in priority order):
  - `empty`: Drive is online and empty.
  - `loading`: Drive is online and is loading a cartridge.
  - `unloading`: Drive is online and unloading writing to the currently loaded cartridge.
  - `reading`: Drive is online and reading from the currently loaded cartridge.
  - `writing`: Drive is online and writing to the currently loaded cartridge.
  - `locating`: Drive is online and doing a seek action on the currently loaded cartridge.
  - `rewinding`: Drive is online and rewinding the currently loaded cartridge.
  - `erasing`: Drive is online and erasing the currently loaded cartridge.
  - `formatting`: Drive is online and formatting the currently loaded cartridge.
  - `calibrating`: Drive is online and calibrating.
  - `ready`: Drive is online, shows no activity, and has a cartridge mounted. This value normally matches the drive's SCSI ready condition (but not always).
  - `unloaded`: Drive is online, shows no activity, and has no cartridge mounted.
- `mtm`: A string that represents the machine type and model of the tape drive (for example, 3588-F6C).

- Bar code: A string that represents the 11S bar code label found on the tape drive. This value is a 22-character string.
- interface: A string that represents the type of ports this tape drive includes. Values include:
  - fibreChannel
  - iSCSI
- interfaceMode: A string that represents the interface mode that is supported by the ports this tape drive includes. Values include: "single-mode" for single-mode fiber and "multi-mode" for multi-mode fiber. The value is null if the ports that are included have an iSCSI interface.
- logicalLibrary: A string that represents the name of the logical library this tape drive is assigned to or null if the tape drive is unassigned.
- use: A string that represents the assigned use that this tape drive was given within the logical library. Values include:
  - access: Data access drives
  - controlPath: Data access and control path drives
  - verification: Media verification drives
- firmware: A string that represents the firmware version that is installed on the tape drive.
- encryption: The state of encryption on this tape drive. Values include:
  - enabled
  - disabled
- wwnn: A 16-character hex string that represents the worldwide node name of the tape drive.
- elementAddress: An integer that represents the element address of the tape drive. This value is reported to hosts by way of the Read Element Status SCSI command. This value is null if unassigned to a logical library.
- beacon: A string that indicates the state of the drive beacon LED. This beacon can be seen from the rear of the drive. Values include:
  - enabled (LED is flashing)
  - disabled (LED is off)

  Note: If the drive is replaced while the beacon is enabled, the state of the physical LED might not match this attribute. In this case, the beacon should be disabled.

- volser: A string that represents the volume serial number or bar code that uniquely identifies the cartridge that is mounted in the tape drive. For example, SG1122L2 or null if the tape drive does not have a cartridge mounted or the currently loaded cartridge's VOLSER is unknown.

Clean drives

Run a drive clean on the specific drive, as shown in the following example:

```
POST /v1/drives/<location>/clean
```

Path parameters

The location parameter is a string that represents the location of the tape drive. This value also is the unique identifier for the drive. The format is "drive_F<f>C<c>R<r>" where "f" is the frame, "c" is the column, and "r" is the row the drive is installed in.

When the **POST /v1/drives** command runs successfully, the system returns output to your display, as shown in Example 6-8.

Example 6-8 POST /v1/drives/drive_F2C4R4/clean command

```
[root@TRON scripts]#/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'POST /v1/drives/drive_F2C4R4/clean'
[root@TRON scripts]#
```

The command prompt returns without comments. However, if we check the status of the “drive_F2C4R4” before and after the command is run, we see that a cleaning cartridge was mounted and the drive is performing a cleaning cycle.

Before and after the POST command, the status of the drive can be checked by issuing write and read buffer **GET /v1/drives/drive_F2C4R4**.

Before command execution (output truncated for better visibility):

```
“state”: “online”,”location”: “drive_F2C4R4”,”operation”: “empty”,”volser”: null
```

After command execution (output truncated for better visibility):

```
“state”: “cleaning”,”location”: “drive_F2C4R4”,”operation”: null,”volser”:
“CLN481JA”
```

Modifying drive use

Modify the use of the drive to a data access, control path, or verification drive, as shown in the following example:

```
PATCH /v1/drives/<location> {“use”: <”access” | “controlPath” | “verification”>}
```

Path parameters

The following path parameters are available:

- location: A string that represents the location of the tape drive. This value also is the unique identifier for the drive. The format is “drive_F<f>C<c>R<r>” where “f” is the frame, “c” is the column, and “r” is the row in which the drive is installed.
- Use: A string that represents the assigned use this tape drive was given within the logical library. Values include:
  - access: Data access drive
  - controlPath: Data access and control path drive
  - verification: Media verification drive

When the **PATCH /v1/drives** command runs successfully, the system returns output to your display, as shown in Example 6-9.

Example 6-9 PATCH /v1/drives/drive_F2C3R1 {“use”: “controlPath”} command

```
[root@TRON scripts]/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'PATCH /v1/drives/drive_F2C3R1 {"use": "controlPath"}'

[root@TRON scripts]
```

The command prompt returns without comments. However, if we check status of the “drive_F2C3R1” before and after the command is run, we see that use was changed to controlPath usage.

Before and after the **PATCH** command, the status of the drive can be checked by issuing write and read buffer **GET /v1/drives/drive_F2C3R1**.

Before command execution (output truncated for better visibility):

`“location”: “drive_F2C3R1”,”use”: “access”,`

After command execution (output truncated for better visibility):

`“location”: “drive_F2C3R1”,”use”: “controlPath”`

Turning drive beacon on/off

Switches on/off the drive beacon LED at the back of the drive, as shown in the following example:

`PATCH /v1/drives/<location> {“beacon”: <”enabled” | “disabled”>}`

Path parameters

The following path parameters are available:

- `location`: A string that represents the location of the tape drive. This value also is the unique identifier for the drive. The format is “drive_F<f>C<c>R<r>” where “f” is the frame, “c” is the column, and “r” is the row in which the drive is installed.
- `beacon`: A string that indicates the current state of the drive beacon LED. This beacon can be seen from the rear of the drive. Values include: “enabled” (LED is flashing) and “disabled” (LED is off).

> **Note:** If the drive is replaced while the beacon is enabled, the state of the physical LED might not match this attribute. In this case, the beacon should be disabled.

When the **PATCH /v1/drives/drive_F2C3R1 {“beacon”: “enabled”}** command runs successfully, the system returns output to your display, as shown in Example 6-10.

Example 6-10 PATCH /v1/drives/drive_F2C3R1 {“beacon”: “enabled”} command

```
[root@TRON scripts]/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'PATCH /v1/drives/drive_F2C3R1 {"beacon": "enabled"}'

[root@TRON scripts]
```

The command prompt returns without comments. However, if we check status of the “drive_F2C3R1” before and after the command is run, we see that the LED is flashing.

Before and after the **PATCH** command, the status of the drive can be checked by issuing write and read buffer **GET /v1/drives/drive_F2C3R1**.

Before command execution (output truncated for better visibility):

`“beacon”: “disabled”,”location”: “drive_F2C3R1”`

After command execution (output truncated for better visibility)

`“beacon”: “enabled”,”location”: “drive_F2C3R1”`

Get Ethernet ports

Retrieve a list of Ethernet ports or a single Ethernet port resource, as shown in the following example:

```
GET /v1/ethernetPorts
GET /v1/ethernetPorts/<location>
```

Path parameter

The location is a parameter string that represents the unique location of this port. The format of this field is “eithernetPort_F<f>P<p>” with possible values for <p> of “a”, “b”, “imc”, “tssc”, or “service”.

When the **GET /v1/ethernetPorts/<location>** command runs successfully, the system returns output to your display, as shown in Example 6-11 on page 411.

Example 6-11 The GET /v1/ethernetPorts/ethernetPort_F1Pa command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg70<<<'GET
/v1/ethernetPorts/ethernetPort_F1Pa'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

[{"location": "ethernetPort_F1Pa", "macAddress": "40:F2:E9:52:71:85",
"ipv4Address": "9.18.77.45", "ipv4Subnet": "255.255.255.0", "ipv4Gateway":
"9.18.77.1", "ipv4Assignment": "static", "ipv4Primary": "9.0.128.50",
"ipv4Secondary": "9.0.130.50", "ipv6Address": "disabled", "ipv6PrefixLength":
"disabled", "ipv6Gateway": "disabled", "ipv6Primary": "disabled", "ipv6Secondary":
"disabled", "ipv6Link": "disabled", "ipv6DHCP": "disabled", "ipv6StatelessConfig":
"disabled", "ipv6Static": "disabled"}]
```

The following command attributes are available:

- location: A string that represents the unique location of this port. The format of this field is “eithernetPort_F<f>P<p>” with possible values for <p> of “a”, “b”, “imc”, “tssc”, or “service”.
- macAddress: The media access control (MAC) address of this Ethernet port.
- ipv4Address: The IPv4 address of this Ethernet port.
- ipv4Subnet: The IPv4 subnet mask of this Ethernet port.
- ipv4Gateway: The IPv4 gateway address of this Ethernet port.
- ipv4Assignment: The IPv4 address of this Ethernet port. Values include:
  - static
  - dynamic
- ipv4Primary: The IPv4 primary DNS address of this Ethernet port. If the ipv4Assignment is “static”, this value is null.
- ipv4Secondary: The IPv4 secondary DNS address of this Ethernet port. If the ipv4Assignment is “static”, this value is null.
- Ipv6Address: The IPv6 address of this Ethernet port.
- ipv6PrefixLength: The IPv6 prefix length of this Ethernet port.
- ipv6Gateway: The IPv6 gateway address of this Ethernet port.
- ipv6Primary: The IPv6 primary DNS address of this Ethernet port.
- ipv6Secondary: The IPv6 secondary DNS address of this Ethernet port.
- ipv6Link: The IPv6 link address of this Ethernet port.
- ipv6DHCP: The IPv6 DHCP setting of this Ethernet port.

- ipv6StatelessConfig: The IPv6 stateless confide setting of this Ethernet port.
- ipv6Static: The IPv6 static address of this Ethernet port.

Get events

Retrieves a list of events or a single event resource. It includes advanced filtering because many events can exist in a library, as shown in the following example:

```
GET /v1/events
GET /v1/events/<ID>
GET /v1/events?{<“after=time” | “before=time” | “location”>}
```

Path parameter

The ID path parameter is a decimal integer that is the unique identifier of the event in the library.

Optional query parameters

The following optional query parameters are available:

- after: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of “YYYY-MM-DDThh:mm:ss+-hhmm” or “YYYY-MM-DDThh:mm:ss”; for example, “2018-09-17T23:02:00”. If the time zone is omitted, the system’s time zone is used.
- before: A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of “YYYY-MM-DDThh:mm:ss+-hhmm” or “YYYY-MM-DDThh:mm:ss”; for example, “2018-09-17T23:02:00”. If the time zone is omitted, the system’s time zone is used.
- location: A string that represents the hardware component where the event originated. Hardware components include:
  - accessor_A<“a”|“b”> (for example, “accessor_Ab”)
  - gripper_A<“a”|“b”>G<”1”|”2”> (for example, “gripper_AbG1”)
  - column_F<f>C<c> (for example, “column_F1C3”)
  - frame_F<f> (for example, “frame_F1”)
  - fiducial_F<f>C<c>L<”t”|”b”> (for example, “fiducial_F1C3IOuLt”)
  - fiducial_F<f>IO<”u”|”l”>L<”t”|”b”> (for example, “fiducial_F1IOuLb”)
  - drive_F<f>C<c>R<r> (for example, “drive_F1C3R23”)
  - slot_ F<f>C<c>R<r>T<t> (for example, “slot_F1C3R23T0”)
  - ioStation_F<f>IO<”u”|”l”> (for example, “ioStation_F1IOu”)
  - ethernetPort_F<f>P<p> (for example, “ethernetPort_F2Pa”)
  - frameSide_F<f>S<”A”|”B”> (for example, “frameSide_F1Sa”)
  - position_F<f>C<c>R<r> (for example, “position_F1C5R7”)

When the **GET /v1/events<after>** command runs successfully, the system returns output to your display, as shown in Example 6-12.

Example 6-12 GET /v1/events?after=2019-11-12T06:22:15-0700 command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg70<<<'GET
/v1/events/after=2019-11-12T06:22:15-0700'
/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

[{"ID": "8275", "severity": "information", "time": "2019-11-12T06:24:54+0000",
"type": "audit", "location": null, "user": "admin", "description": "A user logged
in to the GUI from 9.145.164.92", "state": null, "errorCode": "0800"}, {"ID":
"8274", "severity": "information", "time": "2019-11-12T06:22:29+0000", "type":
"drive", "location": "position_F1C4R1", "user": "System", "description": "The
```

```
state of drive F1, C4, R1 changed from online to unreachable.", "state": null,
"errorCode": "0701"}]
```

The following command attributes are available:

- `ID`: A decimal integer that is the unique identifier of the event in the library.
- `severity`: A string that represents the severity of the event, which indicates the urgency to which that it should be given. Values include:
  - `error`
  - `warning`
  - `inactiveError`
  - `inactiveWarning`
  - `information`
- `time`: The time that the event occurred. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`”; for example, “`2018-09-17T23:02:00-0700`”.
- `location`: A string that represents the hardware component where the event originated. This value can be null if no hardware was associated with this event. Hardware components include:
  - accessor_A<“a”|“b”>
  - gripper_A<“a”|“b”>G<”1”|”2”>
  - frame_F<f>
  - fiducial_F<f>C<c>L<”t”|”b”>
  - fiducial_F<f>IO<”u”|”l”>L<”t”|”b”>
  - drive_F<f>C<c>R<r>
  - slot_ F<f>C<c>R<r>T<t>
  - ioStation_F<f>IO<”u”|”l”>)
  - ethernetPort_F<f>P<p>
  - frameSide_F<f>S<”A”|”B”>
  - position_F<f>C<c>R<r>
- `user`: A string that shows the user name of the user that issued the action that caused this event. This value can be null if this event did not occur as a result of a user action.
- `description`: A string that describes the event that gives more information about what it represents.
- `state`: A string that represents the state of the event regarding its Call Home status. This value can be null if no Call Home yet occurred. Values include:
  - Detected error <error code>
  - Command failed with error code <error code>
  - Calling home
  - Call home initiated
  - Call home failed
  - Call home not attempted because not configured
  - Assigned PMR <PMR number>. Service action required
  - Service action complete by <description>
- `errorCode`: A 4-digit hex string that represents the internal error code that is used to represent this type of event.

Get Frames

These frames include the general information about the installed frames of the library.

Retrieve a list of the currently installed frames of the library, as shown in the following example:

```
GET /v1/frames
GET /v1/frames/<location>
```

Path parameters

The path parameter `location` is a string that represents the unique location of the frame in the library. The format is "`frame_F<f>`" where `'f'` is the frame position.

When the `GET /v1/frames` command runs successfully, the system returns output to your display, as shown in Example 6-13.

Example 6-13 Get /v1/frames command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg15<<<'GET /v1/frames'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg15
[{"location": "frame_F1", "state": "frontDoorOpen", "type": "Base frame", "mtm":
"3584-L25", "sn": "13FA004", "mediaType": "3592","frontDoor": "closed",
    "frontDoorLastChanged": "2021-10-18T14:18:33+0000", "rearDoor": null,
    "rearDoorLastChanged": null, "sideDoor": "closed", "sideDoorLastChanged":
"2021-05-21T20:24:45+0000","slots": 550, "cartridges": 148, "drives": 4,
"ioStations": 2}, {"location": "frame_F2", "state": "frontDoorOpen", "type":
"Storage-only expansion frame", "mtm": "3584-S24", "sn": "13S0051", "mediaType":
"3592", "frontDoor": "closed", "frontDoorLastChanged": "2021-10-18T14:18:33+0000",
"rearDoor": null, "rearDoorLastChanged": null, "sideDoor": "closed",
"sideDoorLastChanged": "2021-05-21T20:24:45+0000","slots": 1000, "cartridges": 21,
"drives": 0, "ioStations": 0}]
```

The following command attributes are available:

- `location`: A string that represents the unique location of the frame in the library. The format is "`frame_F<f>`" where 'f' is the frame position
- `state`: A string that represents the current state of the frame. Values include the following in priority order:
  - `frontDoorOpenWhileNotAllowed`
  - `frontDoorOpen`
  - `sideDoorOpen`
  - `acUnreachable`
  - `calibrationRequired`
  - `inventoryPending`
  - `normal`
- `type`: A string that represents the type of the frame.
- `mtm`: A string that represents the machine type and model of the frame. The following values are available:
  - Base frame
  - Expansion frame
  - Storage-only expansion frame
- `mtm`: A string that represents the machine type and model of this frame.
- `sn`: A string that represents the unique serial number of the frame.

- `mediaType`: A string that represents the media type that is supported by the tape drive. Supported values include “`3592`” and “`LTO`”.
- `frontDoorString` (open, closed)
- `frontDoorLastChangedString` (time/date)
- `rearDoorString` (open, closed)
- `rearDoorLastChangedString` (time/date)
- `sideDoorString` (open, closed)
- `sideDoorLastChangedString` (time/date)
- `slots`: A string that represents the number of available physical slots in this frame.
- `cartridges`: A string that represents the number of cartridges that are in slots or drives in this frame.
- `drives`: A string that represents the number of LTO drives that are installed in this frame.
- `ioStations`: Represents the number of I/O stations that are installed in this frame.

Get Library

This frame includes the general information about the TS4500 tape library and its settings.

Retrieve the library resource, as shown in the following example:

```
GET /v1/library
```

When the **`GET /v1/library`** command runs successfully, the system returns output to your display, as shown in Example 6-14.

Example 6-14 GET /v1/library command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg70<<<'GET /v1/library'
/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

[{"name": "Anacg4", "status": "driveDegraded", "totalCapacity": "1252",
"licensedCapacity": "900", "totalCartridges": "637", "assignedCartridges": "72",
"firmware": "1.6.0.1-B00", "sn": "78AA469", "time": "2019-11-12T07:13:20+0000",
"location": "on earth", "address": "blue Highway", "city": "Nutbush", "state":
"Melmak", "country": "Wonderland", "contact": "Little Joe", "telephone":
"0118-xxx88199", "secondaryTelephone": "xxxx-725-3", "secureCommunications":
"enabled", "autoEjectCleaningCartridges": "enabled", "elasticCapacity": "enabled",
"activeAccessors": "Accessor A only", "vioStatus": "enabled", "nmaDetection":
"enabled", "capacityUtilThresh": "99.0", "dualAccessorUtilThresh": "98.0"}]
```

The following command attributes are available:

- `name`: A string that represents the name of the tape library.
- `status`: A string that represents the overall health status of the library. Values include:
  - `doorOpenWhileNotAllowed`
  - `doorOpen`
  - `pausing`
  - `paused`
  - `restarting`
  - `inServiceMode`
  - `accessorsUnavailable`
  - `accessorDegraded`
  - `nodeCardDegraded`

- updating
- scanningInventory
- online

- `totalCapacity`: An integer that represents the total physical cartridge capacity of the library.
- `licensedCapacity`: An integer that represents the total licensed cartridge capacity of the library.
- `totalCartridges`: An integer that represents the total number of cartridges in the library slots, I/O stations, drives, and accessors.
- `assignedCartridges`: An integer that represents the total number of cartridges that are assigned to logical libraries in the library.
- `firmware`: A string that represents the firmware level that is installed on the library.
- `sn`: A string that represents the serial number of the library.
- `time`: A string that represents the current date and time set on the library. All dates and times that are returned by the library are in relation to this time. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`"; for example, "`2018-09-17T23:02:00-0700`".
- `location`: A string that shows the user entered physical location of the library. For example, "`Building 9062`". If not provided, this value is null.
- `address`: A string that represents the physical address of the library. If not provided, this value is null.
- `city`: A string that represents the city in which the library is located. If not provided, this value is null.
- `state`: A string that represents the two-character state in which the library is located. For example, "`AZ`" for Arizona. If not provided, this value is null.
- `country`: A string that represents the country in which the library is located. If not provided, this value is null.
- `contact`: A string that represents the name of the primary contact for the tape library. This information is used by IBM Support during repair and service. If not provided, this value is null.
- `telephone`: A string that represents the primary telephone number for the primary contact. This information is used by IBM Support during repair and service. If not provided, this value is null.
- `secondaryTelephone`: A string that represents the secondary telephone number for the primary contact. This information is used by IBM Support during repair and service. If not provided, this value is null.
- `secureCommunications`: A string that represents the current setting for secure communications. If `enabled`, the user is required to log in to the GUI by using a secure HTTPS connection. If `disabled`, an unsecured HTTP connection is allowed.
- `autoEjectCleaningCartridges`: A string that represents the current setting for automatically ejecting cleaning cartridges to the I/O Station. If "`enabled`", cleaning cartridges are auto-ejected from the library when they have no cleans remaining. If "`disabled`", they must be ejected manually.
- `elasticCapacity`: A string that represents the current setting for elastic capacity. Values include:
  - maxCapacity
  - tempOverflow
  - doNotUse

- `activeAccessors`: A string that represents which accessors are currently active. Values include:
  - `dualActive`
  - `accessorAOnly`
  - `accessorBOnly`
- `vioStatus`: A string that represents the current setting for Virtual I/O (VIO), which is "`enabled`" or "`disabled`".
- `nmaDetection`: A string that represents the current setting for NMA detection, which is "`enabled`" or "`disabled`".
- `capacityUtilThresh`: A percentage value that is accurate to one decimal place.
- `dualAccessorUtilThresh`: A percentage value that is accurate to one decimal place.

Reset Library

Run a library reset. Although a REST API response always is created, the SCSI Read Buffer method of querying for that response does not necessarily occur before the library reset begins. After resetting, the library status should be monitored until online, as shown in the following example:

```
POST /v1/library/reset
```

When the **`POST /v1/library/reset`** `command` is issued, no response is issued as the library is reset, as shown in Example 6-15 on page 417.

Example 6-15 POST /v1/library/reset command

```
/u/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg70<<<'POST
/v1/library/reset'

<Command will not return a response as the library is reset>
```

Set Library Time, Date, and Time Zone

Set the library time, date, and time zone, as shown in the following example:

```
PATCH /v1/library {"time": <time>}
```

Request Body parameters

The `time` parameter is the current date and time to be set on the library. All dates and times that are returned by the library are in relation to this time. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`"; for example, "`2018-09-17T23:02:00-0700`".

When the **`PATCH /v1/library {"time": <time>}`** command runs successfully, the system returns output to your display, as shown in Example 6-16.

Example 6-16 PATCH /v1/library {"time": "2019-11-12T10:07:00-0600"} command

```
[root@TRON scripts]/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'PATCH /v1/library {"time": "2019-11-12T10:09:00-0600"}'

[root@TRON scripts]
```

The time changes and a response regarding the command's completion is returned. You can check the GUI for changes by selecting **Settings** → **Library** → **Date and Time**. The date and time is the same as set by the **sg_write_buffer** command with a time zone of -06:00h (see Figure 6-1).

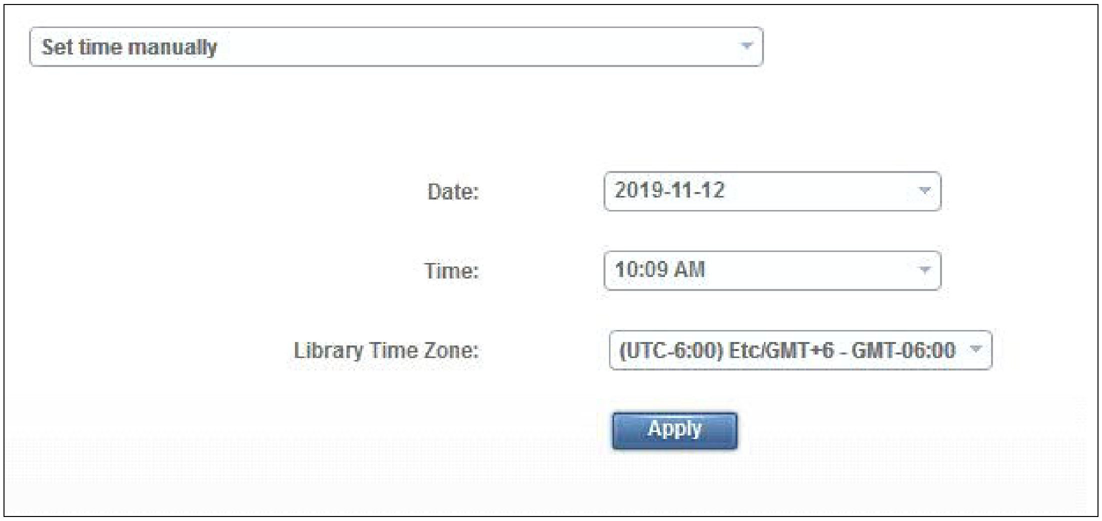

Figure 6-1 Set Library Date and Time

LOGS

LOGS retrieve a list of library logs, create a library log, or export a created library log, as shown in the following example:

```
GET /v1/logs
POST /v1/logs
GET /v1/logs/<filename>
GET /v1/logs/<filename>/export
```

Path parameters

The filename path parameter is the name of the file that is created or in the process of being created. This file name contains the name of the library, the date and time it was taken, and other important service information.

Start creating the logs. When creating the log file, no feedback is shown. With the **GET/v1/logs** command, the generate logs are shown.

Creating a log file is shown in Example 6-17.

Example 6-17 POST v1/logs command

```
[root@TRON scripts]# /usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg15<<<'POST /v1/logs'
```

Get logs

Retrieve a list of logs that exist on the library. Some logs on this list might still be in the process of being created and can be monitored by using the progress attribute.

With the **GET /v1/logs** command, all generated log files can be displayed.

When the **GET /v1/logs** command runs successfully, the system returns output to your display, as shown in Example 6-18.

Example 6-18 Get /v1/logs command

```
GET /v1/logs
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 45992

[
{
"filename": "log1.zip",
"state": "complete",
"startTime": "2020-02-11T11:13:44-0700",
"lastUpdateTime": "2020-02-11T11:15:55-0700",
"percentComplete": 100
}, {
"filename": "log2.zip",
"state": "inProgress",
"startTime": "2020-03-16T17:14:41-0700",
"lastUpdateTime": "2020-03-16T17:15:55-0700",
"percentComplete": 61
```

The following command attributes are available:

- `filename`: The name of the file that is created or in the process of being created. This file name contains the name of the library, the date and time it was taken, and other important service information.
- `state`: A string that shows the completion status of the event. Values include:
  - `inProgress`: The task is in progress. If supported, the **"**`percentComplete`**"** attribute shows how far along the task is and the **"**`duration`**"** attribute shows how long the task was running.
  - `completed`: The task completed successfully. The file can be offloaded by using the **GET /v1/logs/<ID>/export** command.
  - `failed`: The task failed. Retry the attempt. If it continues to fail, contact IBM Support.
- `stratum`: The time at which the log creation task was started. Time format is in ISO 8601 format of **"**`YYYY-MM-DDThh:mm:ss+-hhmm`**"**; for example, **"**`2020-03-17T23:02:00-0700`**"**.
- l`astUpdatedTime`: The time this task's state was last updated. Time format is in ISO 8601 format of **"**`YYYY-MM-DDThh:mm:ss+-hmm`**"**; for example, **"**`2020-03-18T22:01percentComplete`**"**.
- `percentComplete`: A percentage value that is accurate to one decimal place and shows how close the task is to being complete.

Export logs

The **`GET /v1/logs/<filename>/export`** command is used.

Export a created log. This command exports the log file of the specified filename. The output is shown is Example 6-19.

Example 6-19 GET /v1/logs/<filename>/export command

```
[root@TRON scripts]# ./GetLogs.sh sg15 TS4500_LOG_FA004_20200317081743.zip
Getting piece 1
Done with piece 1
Length = 0xffffff
Getting piece 2
Done with piece 2
Length = 0xffffff
Getting piece 3
```

The log file is locally saved on the host where the **`Get /v1/logs/<filename>/export`** command is run.

Note: Our example script, Example 6-19, was used to get the log file. During the transfer from the library to the local host the data is grouped in 1Mb pieces

Get Node Cards

Retrieve a list of node cards or a single node card resource, as shown in the following example:

```
GET /v1/nodeCards
GET /v1/nodeCards/<ID>
POST /v1/nodeCards/<ID>/reset
```

Path parameter

The path parameter `ID` is a decimal integer that is the unique identifier of the node card in the library.

When the **`GET /v1/nodeCards/<ID>`** command runs successfully, the system returns output to your display, as shown in Example 6-20 on page 420.

Example 6-20 GET /v1/nodeCards/17 command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'GET
/v1/nodeCards/17'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

[{"ID": "17", "type": "MDA", "location": "accessor_Aa", "state": "online",
"partNum": "38L7590", "sn": "Y010MY95Z61X", "bar code": "11S38L7590Y010MY95Z61X",
"ec": "0", "firmware": "0068", "cfBarcode": null, "cfPartNum": null, "cfVendor":
null, "reportingLCC": null, "primaryLCC": null}]
```

The following command attributes are available:

- `ID`: A decimal integer that is the unique identifier of the node card in the library.
- `type`: A string that represents the physical type of the node card. Values include:
  - `MDA`: Motor driver assembly
  - `ACC`: Accessor control card
  - `LCC`: Library control card
- `location`: A string that represents the hardware component where this node card is located. For the MDA card and ACC, this value identifies the accessor in the format **"**`accessor_A<a>`**"**. For the LCC card, this value identifies the frame in the format **"**`frame_F<f>`**"**.
- `state`: A string that represents the current state of the node card. Values include:
  - `inServiceMode`
  - `unreachable`
  - `noEthernet`
  - `noCAN`
  - `online`
- `partNum`: A string that represents the part number of the node card.
- `sn`: A string that represents the serial number of the node card.
- `Bar code`: A string that represents the exact bar code label found on the node card. This string is 22-characters.
- `ec`: An integer representing the engineering change (EC) number of the node card.
- `firmware`: A string that represents the firmware level of the node card.
- `cfBarcode`: A string that represents the value read from the serial number bar code of the compact flash card that is installed on the LCC. If this card is not an LCC card, this value is null. The format of the bar code matches that read off the compact flash card. Values include:
  - `SMART`
  - `Virtium`
- `cfPartNum`: A string that represents the part number of the compact flash card that is installed on the LCC. If this card is not an LCC or the CF card does not support reporting this value, the value is null.
- `cfVendo`: A string that represents the name of the vendor of the compact flash card that is installed on the LCC. If this card is not an LCC card, this value is null. The other two supported values include:
  - `SMART`
  - `Virtium`
- `reportingLCC`: A string that represents whether this LCC is reporting the information for this **`GET`** command. This command returns `yes` if this LCC is reporting the status, `no` if this LCC's status is being reported by another LCC, or `null` if this card is not an LCC card.
- `primaryLCC`: A string that represents whether this LCC is acting as the primary LCC or the secondary LCC. This command returns `yes` if this LCC is the primary, `no` if this LCC is the secondary, or `null` if this card is not an LCC card.

Power supplies

The TS4500 tape library connects to AC power in many ways. The library is powered by two AC to DC power supplies in the Lx5 frame. You can also have power supplies that are installed in a Dx5 frame. A **GET** command can be used to list attributes of the power supplies in all frames, as shown in the following example:

```
GET /v1/powerSupplies
```

When the **GET /v1/powerSupplies** command runs successfully, the system returns output to your display, as shown in Example 6-22.

Example 6-21 GET /v1/powerSupplies

```
[root@TRON scripts]# /usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg8<<<'GET /v1/powerSupplies'
[root@tron ITDT]# /usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000
--raw /dev/sg8
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 216

[{"location": "powerSupply_F1PSa", "state": "online"}, {"location":
"powerSupply_F1PSb", "state": "online"}, {"location": "powerSupply_F2PSa",
"state": "online"}, {"location": "powerSupply_F2PSb", "state":
"online"}][root@tron ITDT]#

[root@TRON scripts]
```

Attributes

The following attributes are available:

- Location: The string location of the power that is supplying the format:

  "`powerSupply_F<frame>PS<power supply A or B>`"

  For example: "`powerSupply_F1PSa`".

- State: The current operation mode of the power supply as last read by the LCC node card (can be obsolete if LCC node card is unreachable). The following values are included:
  - `failed`

    The power supply is showing a failure or is unreachable.Verify that the power supply is plugged in and cabled correctly from the rear of the frame. If this check does not resolve the issue, contact IBM for support.
  - `online`

    The power supply is plugged in and supplying power to the library.
  - `inServiceMode` *(future)*

    The power supply is powered off in preparation for service.

Reset Node Card

Run a node card reset on the specific node card. If the node card is the LCC that runs this command, a REST API response always is created; however, the SCSI Read Buffer method of querying for that response does not necessarily occur before the node card reset begins. However, the buffer is non-volatile and query-able after the reset is complete (see the following example).

```
POST /v1/nodeCards/<ID>/reset
```

When the **`POST /v1/nodeCards/<ID>/reset`** command runs successfully, the system returns output to your display, as shown in Example 6-22.

Example 6-22 POST /v1/nodeCards/65/reset command

```
[root@TRON scripts]# /usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'POST /v1/nodeCards/65/reset'

[root@TRON scripts]
```

Node 65 is the LCC node card in frame 1. A reset of this node card causes the LCC to restart and post an event, as shown in the following examples:

- `Error Code: 2384`
- `Description: An LCC card was not responding over Ethernet`

The TS4500 has the following node cards installed:

- NodeIdx 17 (x'11') XYC Acc A, NodeIdx 18 (x'12') XYC Acc B
- NodeIdx 49 (x'31') ACC Acc A, NodeIdx 50 (x'32') ACC Acc B
- NodeIdx 65 (x'41') LCC Frame 1
- NodeIdx 66 (x'42') LCC Frame 2
- NodeIdx 66 (x'42') LCC Frame 3, and so on

Report

The reports contain usage history and other data for resources in the library. A report is available for the following components:

- Drives
- Library
- Accessors

Get drives report

The drive report contains usage history of the drive including mounts, cleans, host I/O, compression rate, error rates, and environmental data. These reports are taken at 1-hour intervals and are kept for up to 1 year. However, you can only offload 1 week of data at a time per **`GET`** command to ensure that the query response size does not get too large. If you do not include query parameters to limit the time frame, all reports for the last week are returned, as shown in the following example:

```
GET/v1/reports/drives
GET/v1/reports/drives?{<“after=time” | “before=time” | “location”>}
```

Optional query parameters

The following optional parameters are available:

- after: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-17T13:02:00". If the time zone is omitted, the current systems current time zone is used.
- before: A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2020-03-17T13:02:00". If the time zone is omitted the current systems current time zone is used.

When the **GET /v1/reports/drives** command runs successfully, the system returns output to your display, as shown in Example 6-23.

Example 6-23 Get /v1/report/drives command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg15<<<'GET
/v1/reports/drives'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg15
[{"location": "drive_F1C4R1", "sn": "0000078PG12E", "time":
"2020-03-16T12:05:00-0600", "duration": 3600, "mounts": 0, "cleans": 0,
"dataReadByHosts": 0, "dataWrittenByHosts": 0, "dataWrittenToCartridges": 0,
"errorsCorrectedRead": 0, "errorsCorrectedWrite": 0, "errorsUncorrected": 0,
"temperatureAverage": 22.0, "temperatureMin": 22.0, "temperatureMax": 22.0,
"humidityAverage": 53.0, "humidityMin": 22.0, "humidityMax": 53.0},
```

Note: Only one drive is displayed in Example 6-23. For example, when a library has 72 drives, 72 drives are displayed.

The following command attributes are available:

- location: A string that represents the unique location of the tape drive. This value also is the unique identifier for the drive. The format is **"**drive_F<f>C<c>R<r>**"** where 'f' is the frame, 'c' is the column, and 'r' is the row in which the drive is installed.
- sn: A string that represents the serial number of the drive. This value also can be used as a unique identifier for the tape drive.
- time: The time at which this data entry was recorded. Time format is in ISO 8601 format of **"**YYYY-MM-DDThh:mm:ss+-hhmm**"**; for example, **"**2020-03-17T13:02:00-0700**"**.
- duration: The number of seconds this data entry represents in seconds. This value is approximately 3600 seconds, or 1 hour, but can differ slightly because of execution time. The 'time' attribute occurs at the end of this duration.
- mounts: An integer that represents the number of mounts this is performed per hour on this drive over the duration of this measurement.
- cleans: An integer that represents the number of cleans are performed per hour on this drive over the duration of this measurement.
- dataReadByHost: An integer that represents the number of MB per hour of uncompressed data read by the host from the drive over the duration of this measurement.

- `dataWrittenToCartridges`: An integer that represents the number of MB per hour of compressed data that is written to the cartridges that are mounted in the drive over the duration of this measurement. The average compression ratio can be found by dividing `dataWritenByHosts` by `dataWrittenToCartridge`.
- `errorsCorrectedRead`: An integer that represents the number of corrected read errors per hour on the drive over the duration of this measurement.
- `errorsCorrectedWrite`: An integer that represents the number of corrected write errors per hour on the drive over the duration of this measurement.
- `errorsUncorrected`: An integer that represents the number of uncorrected errors per hour on the drive over the duration of this measurement.
- `temperatureAverage`: A floating-point number that is accurate to one decimal place and represents the average temperature in Celsius of the drive over the duration of this measurement.
- `temperatureMin`: A floating-point number that is accurate to one decimal place and represents the minimum temperature in Celsius that is reached in the drive over the duration of this measurement.
- `temperatureMax`: A floating-point number that is accurate to one decimal place and represents the minimum temperature in Celsius that is reached in the drive over the duration of this measurement.
- `humidityAverage`: A floating-point number that is accurate to one decimal place and represents the average percent humidity in the drive over the duration of this measurement. This value is null if the drive does not support reporting humidity measurements.
- `humidityMin`: A floating-point number that is accurate to one decimal place and represents the minimum (driest) percent humidity that is reached in the drive over the duration of this measurement. This value is null if the drive does not support reporting humidity measurements.
- `humidityMax`: A floating-point number that is accurate to one decimal place and represents the highest (wettest) percent humidity that is reached in the drive over the duration of this measurement. This value is null if the drive does not support reporting humidity measurements.

Get library report

The library report contains usage history of the library including mounts, inserts, ejects, moves, host I/O, and environmental data. These reports are taken on intervals of 1 hour and kept for 1 year. However, you can only offload 1 week of data at a time per **`GET`** command to ensure that the query response size does not get too large. If you do not include query parameters to limit the time frame, all reports for the last week are returned, as shown in the following example:

```
GET/v1/reports/library
GET/v1/reports/library?{<“after=time” | “before=time” | “location”>}
```

Optional query parameters

The following optional parameters are available:

- `after`: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`” or “`YYYY-MM-DDThh:mm:ss`”; for example, “`2020-03-17T13:02:00`”. If the time zone is omitted the current systems current time zone is used.

- `before`: A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`"; for example, "`2020-03-17T13:02:00`". If the time zone is omitted, the current systems current time zone is used.

When the **`GET /v1/reports/library`** command runs successfully, the system returns output to your display, as shown in Example 6-24.

Example 6-24 GET /v1/reports/library command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg15<<<'GET
/v1/reports/library'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg15
[{"time": "2020-03-17T07:05:03-0600", "duration": 3600, "mounts": 430, "imports":
111, "exports": 185, "moves": 1238, "dataReadByHosts": 12934568,
"dataWrittenByHosts": 65973285, "dataWrittenToCartridges": 872320954,
"temperatureAverage": 24.0, "temperatureMin": 21.0, "temperatureMax": 26.0,
"humidityAverage": 45.0, "humidityMin": 40.0, "humidityMax": 49.0},
```

The following command attributes are available:

- `time`: The time at which this data entry was recorded. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`"; for example, "`2018-09-17T23:02:00-0700`".
- `duration`: The number of seconds this data entry represents in seconds. This value is approximately 3600 seconds, or 1 hour, but can differ slightly because of execution time. The `'time'` attribute occurs at the end of this duration.
- `mounts`: An integer that represents the number of mounts that are performed per hour on all drives in this library over the duration of this measurement.
- `imports`: An integer that represents the number of cartridges that are added to a library per hour over the duration of this measurement. A cartridge import is not complete until the host sends a SCSI move media command to the library for this cartridge or it was manually assigned to a Logical Library.
- `exports`: An integer that represents the number of cartridges that are moved to the I/O station from a slot per hour over the duration of this measurement. A cartridge export is complete after the cartridge is physically moved to the I/O station.
- `moves`: An integer that represents the number of cartridge movement actions (including the get and put by the gripper) per hour that are run by all accessors in the library over the duration of this measurement.
- `dataReadByHosts`: An integer that represents the number of MB per hour of uncompressed data read by the host from all drives in the library over the duration of this measurement.
- `dataWrittenByHosts`: An integer that represents the number of MB per hour of uncompressed data that is written by the host to all drives in the library over the duration of this measurement.
- `dataWrittenToCartridges`: An integer that represents the number of MB per hour of compressed data that is written to the cartridges by all drives in the library over the duration of this measurement. The average compression ratio can be found by dividing `dataWritenByHosts` by `dataWrittenToCartridge`.
- `temperatureAverage`: A floating-point that is accurate to one decimal place number that represents the average temperature in Celsius in all drives in the library over the duration of this measurement.

- `temperatureMin`: A floating-point that is accurate to one decimal place number that represents the minimum temperature in Celsius that is reached in all drives in the library over the duration of this measurement.
- `temperatureMax`: A floating-point number that is accurate to one decimal place that represents the minimum temperature in Celsius that is reached that is accurate in all drives in the library over the duration of this measurement.
- `humidityAverage`: A floating-point number that is accurate to one decimal place that represents the average percent humidity that is reached in all drives in the library over the duration of this measurement.
- `humidityMin` A floating-point number that is accurate to one decimal place that represents the minimum (driest) percent humidity that is reached in all drives in the library over the duration of this measurement.
- `humidityMax`: A floating-point number that is accurate to one decimal place that is represents the highest (that is, wettest) percent humidity that is reached in all drives in the library over the duration of this measurement.

Get Accessor Report

The accessor report contains usage history of the accessor, including pivots, scans, distance traveled, and puts and gets that are done by the two grippers. These reports are taken at 1-hour intervals and kept for 1 year. However, you can offload only 1 week of data at a time per **`GET`** command to ensure that the query response size does not get too large. If you do not include query parameters to limit the time frame, all reports for the last week are returned (see the following example).

```
GET/v1/reports/accessors
GET/v1/reports/accessors?{<"after=time" | "before=time" | "location">}
```

Option query parameters

The following optional query parameters are available:

- `after`: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`"; for example, "`2020-03-17T13:02:00`". If the time zone is omitted, the current system's time zone is used.
- `before`: A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`"; for example, "`2020-03-17T13:02:00`". If the time zone is omitted, the current system's time zone is used.

When the **`GET/v1/reports/accessors`** command runs successfully, the system returns output to your display, as shown in Example 6-25.

Example 6-25 GET /v1/reports/accessors

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg15<<<'GET
/v1/reports/accessors'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg15
[{"location": "accessor_Aa", "time": "2020-03-17T09:05:00-0600", "duration": 3600,
"pivots": 4458, "barCodeScans": 23423, "travelX": 42442, "travelY": 2423440,
"getsGripper1": 342340, "putsGripper1": 2343240, "getsGripper2": 424230,
"putsGripper2": 2342340},]
```

The following command attributes are available:

- `location`: A string that describes the unique identifier for the accessor with a value of "`accessor_Aa`" or "`accessor_Ab`".
- `time`: The time at which this data entry was recorded. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`". For example, "2020-03-18T13:02:00-0700".
- `duration`: The number of seconds this data entry represents in seconds. This value is approximately 3600 seconds, or 1 hour, but might differ slightly because of the run time. The `'time'` attribute occurs at the end of this duration.
- `pivots`: An integer that represents the number of pivots per hour that are performed by this accessor over the duration of this measurement.
- `barCodeScans`: An integer that represents the number of bar code scans per hour that are performed by this accessor over the duration of this measurement.
- `travelX`: An integer that represents the number meters per hour of movement in the X (horizontal) direction that are performed by this accessor over the duration of this measurement.
- `travelY`: An integer that represents the number meters per hour of movement in the Y (vertical) direction that are performed by this accessor over the duration of this measurement.
- `getsGripper1`: An integer that represents the number cartridge get actions per hour that are performed by gripper 1 of this accessor over the duration of this measurement.
- `putsGripper1`: An integer that represents the number cartridge put actions per hour that are performed by gripper 1 of this accessor over the duration of this measurement.
- `getsGripper2`: An integer that represents the number cartridge get actions per hour that performed by gripper 2 of this accessor over the duration of this measurement.
- `putsGripper2`: An integer that represents the number cartridge put actions per hour that are performed by gripper 2 of this accessor over the duration of this measurement.

Slots

This command returns all storage slots in the system. Default sorting is by location in numerical order that is based on Frame, Column, and Row:

```
GET /v1/slots
```

When the **`GET /v1/slots`** command runs successfully, the system returns output to your display, as shown in Example 6-26. The output is truncated because all slots in the library were listed.

Example 6-26 GET /v1/slots

```
[root@tron ITDT]# /usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg8<<<'GET /v1/slots'
[root@tron ITDT]# /usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000
--raw /dev/sg8
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 82668

[{"location": "slot_F1C3R1", "state": "inServiceMode", "contents": [null, null,
null, null], "puts": 59, "putRetries": 0, "getRetries": 12, "tiers": 4},
{"location": "slot_F1C3R2", "state": "inServiceMode", "contents": [null, null,
null, null], "puts": 46, "putRetries": 0, "getRetries": 9, "tiers": 4},
{"location": "slot_F1C3R3", "state": "inServiceMode", "contents": [null, null,
```

```
null, null], "puts": 38, "putRetries": 0, "getRetries": 3, "tiers": 4},
{"location": "slot_F1C3R4", "state": "inServiceMode",

[root@TRON scripts]
```

Attributes

The following attributes are available:

- Location

  The string location of the slot in the library, not including the tier; for example, `slot_F3C4R8`.

  A fully populated deep slot is shown in the following example:

  ```
  "location": "slot_F1C9R1", "state": "normal", "contents": [ "JJL051JJ",
  "CD0347JA",
    "GH0141JB", "CE 211JJ" ], "puts": 93, "putRetries": 0, "getRetries": 16, "tiers": 4
  ```

- State

  The current state of the slot, which includes the following values:

  – `inServiceMode`

    It is not selected as a cartridge destination; however, cartridges can be moved from the slot.

  – `unusable` (*future*)

    Not in use because it is not available for media move operations because the elastic capacity is being disabled.

  – `inaccessible` (*future*)

    The slot is not accessible by an accessor. This issue occurs because an accessor in the `noMovementAllowed` state prevents access to this slot, or all accessors are in a state that is unable to access.

  – `normal`

    Slot is ready for use.

- Contents

  An array of strings that shows the Volsers of the cartridges that are contained in the slot. They are listed in order of lowest numbered tier to highest numbered tier. Any empty tier is listed as null.

  Note: This attribute is reported as a single-item array for a tier 0 (single-deep, door-side) slot.

- Puts

  An integer that shows the number of times a gripper inserted a cartridge into this slot. This attribute includes all put actions, including shuffle and unshuffle moves.

- `putRetries`

  An integer that shows the number of times a retry was required while a gripper was inserting a cartridge into this slot.

- `getRetries`

  An integer that shows the number of times a retry was required while a gripper was retrieving a cartridge from this slot.

- Tiers

  An integer that represents the number of tiers within this slot.

Tasks

Retrieve a list of tasks or a single task resource:

```
GET /v1/tasks
GET /v1/tasks/<ID>
```

When the **`GET /v1/tasks/<ID>`** command runs successfully, the system returns output to your display, as shown in Example 6-27.

Example 6-27 GET /v1/tasks/140 command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'GET /v1/tasks/140'
/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70
[{"ID": 140, "type": "updateLibraryFirmware", "location": "library", "state":
"completed", "startTime": "2019-08-15T09:21:32-0600", "lastUpdateTime":
"2019-08-15T09:34:30-0600", "percentComplete": 100, "duration": 778, "user":
"admin", "description": "Library Code Update", "volser": null}]
```

Path parameters

The following path parameters are available:

- `ID`: A decimal integer that is the unique identifier of the task in the library.
- `type`: A string that represents the task type to run. Values can include:
  - `inventoryTier0and1`
  - `inventoryAllTiers`
  - `calibrateLibrary`
  - `calibrateFrame`
  - `calibrateAccessor`
  - `testDrive`
  - `startDriveService`
  - `completeDriveService`
  - `startAccesorService`
  - `completeAccesorService`
- `location`: A string that represents the hardware component what the task is affecting. This value can be null if no hardware was associated with this task. Hardware components include:
  - `accessor_A`<“a”|”b”> (for example, “`accessor_Ab`”)
  - `gripper_A`<“a”|”b”>`G`<“1”|”2”> (for example, “`gripper_AbG1`”)
  - `frame-up`> (for example, “`frame_F1`”)
  - `drive_F<f>C<c>R<r>` (for example, "`drive_F1C3R23`")
  - `ethernetPort_F<f>P<p>` (for example, “`ethernetPort_F2Pa`”)
- `state`: A string that shows the completion status of the task. Values include:
  - `inProgress`: The task is in progress. If supported, the “`percentComplete`” Task attribute shows how far along the task is and the `duration` attribute shows how long the task was running.
  - `completed`: The task completed successfully.
  - `failed`: The task failed. To find the reason for the failure, check the event list.
  - `aborted`: An LCC failover caused the task to be canceled automatically.
  - `canceled`: A user manually canceled the task or the library powered off.

- `startTime`: The time the task was started. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`"; for example, "`2018-09-17T23:02:00-0700`".
- `lastUpdatedTime`: The time this task's state was last updated. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`"; for example, "`2018-09-17T23:02:00-0700`".
- `percentComplete`: A percentage value that is accurate to one decimal place that shows how close the task is to being complete. This value is an estimate of the time the task takes and is *not* a guarantee. This value is null for tasks that do not report this information.
- `duration`: An integer that represents the number of seconds this task was running so far.
- `user`: A string that shows the name of the user whose action started the task.
- `volser`: For event types of "`verifyMedia`", this field shows the VOLSER being verified. For other event types, this value is null.

Inventory Tier 0 and 1

Run an inventory scan on the library. This scan starts a long-running task in the library that is visible from the GUI, as shown in the following example:

```
POST /v1/tasks {"type": "inventoryTier0and1", "location": <"library" |
"frame_F<f>">}
```

Request body parameters

The following request body parameters are available:

- `type`: A string that represents the task type to run; in this `case` "`inventoryTier0and1`".
- `location`: This value can be "`library`" for all frames, or the location of the frame to inventory; for example, "`frame_F4`".

When the **POST /v1/tasks {"type": "inventoryTier0and1", "location": "frame_F<f>">}** command runs successfully, the system returns output to your display, as shown in Example 6-28.

Example 6-28 POST /v1/tasks/ {"type": "inventoryTier0and1", "location": "frame_F1"} command

```
[root@TRON scripts]# /usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'POST /v1/tasks/ {"type": "inventoryTier0and1", "location":
"frame_F1"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

{"ID": "331"}
```

Task ID 331 was generated and you can see that inventory started for Frame 1 at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/331** command.

Inventory all tiers

Run an inventory scan on the library. This scan starts a long-running task in the library that is visible from the GUI, as shown in the following example:

```
POST /v1/tasks [{"type": "inventoryAllTiers", "location": <"library" |
"frame_F<f>">}]
```

Request body parameters

The following request body parameters are available:

- type: A string that represents the task type to run; in this case "inventoryTier0and1".
- location: This value can be library for all frames, or the location of the frame to inventory; for example, "frame_F4".

When the **POST /v1/tasks [{"type": "inventoryAllTiers", "location": "frame_F<f>">}]** command runs successfully, the system returns output to your display, as shown in Example 6-29.

Example 6-29 POST /v1/tasks/ {"type": "inventoryAllTiers", "location": "frame_F2"} command

```
[root@TRON scripts]# /usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw
/dev/sg70<<<'POST /v1/tasks/ {"type": "inventoryAllTiers", "location":
"frame_F2"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

{"ID": "333"}
```

Task ID 333 was generated and you can see that inventory with audit started for Frame 2 at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/333** command.

Start service on drive

Start a service action on the specific drive. This action puts the drive in the "in service" state and starts a long-running task in the library that is visible from the GUI:

```
POST /v1/tasks [{"type": "startDriveService", "location": "drive_F<f>C<c>R<r>"}]
```

Request body parameters

The following list describes the parameters:

- type: A string that represents the task type to run. In this case "startDriveService".
- location: A string that represents the location of the tape drive. This is also the unique identifier for the drive. The format is "drive_F<f>C<c>R<r>" where "f" is the frame, "c" is the column, and "r" is the row the drive is installed in.

When the **POST /v1/tasks [{"type": "startDriveService", "location": "drive_F<f>C<c>R<r>"}]** command runs successfully, the system returns output to your display, as shown in Example 6-30.

Example 6-30 POST /v1/tasks [{"type": "startDriveService", "location": "drive_F1C2R2}] command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'POST /v1/tasks/
{"type": "startDriveService", "location": "drive_F1C2R2"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

{"ID": "335"}
```

Task ID 335 was generated and you can see that "Drive FRU Start" task was created for drive F1 C2 R2 at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/335** command. The drive state is offline.

Complete service on drive

Complete a service action on the specified drive. This process takes the drive out of the “in service” state and starts a long-running task in the library that is visible from the GU:

POST /v1/tasks [{“type”: ”completeDriveService”, “location”: ”drive_F<f>C<c>R<r>”}]

Request body parameters

The following request body parameters are available:

- type: A string that represents the task type to run; in this case, “completeDriveService”.
- location: A string that represents the location of the tape drive. This value is also the unique identifier for the drive. The format is “drive_F<f>C<c>R<r>” where “f” is the frame, “c” is the column, and “r” is the row in which the drive is installed.

When the **POST /v1/tasks [{“type”: ”completeDriveService”, “location”: ”drive_F<f>C<c>R<r>”}] com**mand runs successfully, the system returns output to your display, as shown in Example 6-31.

Example 6-31 POST /v1/tasks [{“type”: ”completeDriveService”, “location”: ”drive_F1C2R2”}] command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'POST /v1/tasks/
{"type": "completeDriveService., "location": "drive_F1C2R2"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

{"ID": "337"}
```

Task ID 337 was generated and you can see that “Drive FRU Finish” task was created for drive F1 C2 R2 at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/337** command.

Calibrate accessor

Run a calibration on the specified accessor. This process starts a long-running task in the library that is visible from the GUI:

POST /v1/tasks [{“type”: ”calibrateAccessor”, “location”: ”accessor_A<”a”|”b”>”}]

Request body parameters

The following parameters are available:

- type: A string that represents the task type to run; in this case, “calibrateAccessor”.
- location: A string that describes the unique identifier for the accessor with a format of “accessor_A<“a”|“b”>”.

When the **POST /v1/tasks [{“type”: ”calibrateAccessor”, “location”: ”accessor_A<”a”|”b”>”}]** command runs successfully, the system returns output to your display, as shown in Example 6-32.

Example 6-32 POST /v1/tasks/ {“type”: “calibrate Accessor., “location”: “accessory”} command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'POST /v1/tasks/
{"type": "calibrateAccessor., "location": "accessor_Aa"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70
{"ID": "339"}
```

Task ID 339 was generated and you can see that “Calibrate Accessor” task was created for Accessor A at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/339** command.

Start service on accessor

Start a service action on the specified accessor. This process puts the accessor in the “in service” state, which starts a long-running task in the library that is visible from the GUI:

```
POST /v1/tasks [{“type”: ”startAccessorService”, “location”:
”accessor_A<”a”|”b”>”}]
```

Request body parameters

The following request body parameters are available:

- type: A string that represents the task type to run; in this case, “start Accessor Service”.
- location: A string that describes the unique identifier for the accessor with a format of “accessor_A<“a”|“b”>”.

When the **POST /v1/tasks [{“type”: ”startAccessorService”, “location”: ”accessor_A<”a”|”b”>”}]** command runs successfully, the system returns output to your display, as shown in Example 6-33.

Example 6-33 POST /v1/tasks/ {“type”: "startAccessorService., “location”: "accessor_Aa"} command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'POST /v1/tasks/
{"type": "startAccessorService., "location": "accessor_Aa"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

{"ID": "341"}
```

Task ID 341 was generated and you can see that “Accessor FRU Start” task was created for Accessor A at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/341** command.

Complete service on accessor

Complete a service action on the specified accessor. This process takes the accessor out of the “in service” state and starts a long-running task in the library that is visible from the GUI:

```
POST /v1/tasks [{“type”:”completeAccessorService”, “location”:”
accessor_A<”a”|”b”>”}]
```

Request body parameters

The following request body parameters are available:

- type: A string that represents the task type to run; in this case, “completeAccessorService”.
- location: A string that describes the unique identifier for the accessor with a format of “accessor_A<“a”|“b”>”.

When the **POST /v1/tasks [{“type”: ”completeAccessorService”, “location”: ”accessor_A<”a”|”b”>”}]** command runs successfully, the system returns output to your display, as shown in Example 6-34.

Example 6-34 POST /v1/tasks/ {“type”: “completeAccessorService, “location”: “accessor_Aa”} command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'POST /v1/tasks/
{"type": "completeAccessorService, "location": "accessor_Aa"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

{"ID": "343"}
```

Task ID 343 was generated and you can see that “`Accessor FRU Finish`” task was created for Accessor A at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/343** command.

The following steps are performed by the accessor during the “`Complete Service`” operation:

1. FRU Complete ACC-A - Rezero Request
2. FRU Complete ACC-A - Calibrate Accessor
3. FRU Complete ACC-A - Calibrate ISB (Internal Service Bay)
4. FRU Complete ACC-A - Verify Library
5. FRU Complete ACC-A - Issue Final Rezero
6. FRU Complete ACC-A - Done

Test the drive

Start a test drive operation by using a diagnostic cartridge. This process starts a long-running task in the library that is visible from the GUI:

```
POST /v1/tasks [{“type”: “testDrive”, “location”: “drive_F<f>C<c>R<r>”}]
```

Request body parameters

The following parameters are available:

- `type`: A string that represents the task type to run; in this case, “`testDrive`”.
- `location`: A string that represents the location of the tape drive. This value is also the unique identifier for the drive. The format is “`drive_F<f>C<c>R<r>`” where “`f`” is the frame, “`c`” is the column, and “`r`” is the row in which the drive is installed.

When the **POST /v1/tasks [{“type”: “testDrive”, “location”: “drive_F<f>C<c>R<r>”}]** command runs successfully, the system returns output to your display, as shown in Example 6-35.

Example 6-35 POST /v1/tasks/ {“type”: “testDrive, “location”: “driveF1C2R2”} command

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw /dev/sg70<<<'POST /v1/tasks/
{"type": "testDrive", "location": "driveF1C2R2"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg70

{"ID": "345"}
```

Task ID 345 was generated and you can see that “`Test Drive`” task was created for drive F1 C2 R2 at the web GUI. You can also determine the status and progress by using the **GET /v1/tasks/345** command.

Work items

Moves a cartridge to the specified drive:

```
POST /v1/workItems [{"type": "moveToDrive", "cartridge":
<volser>,"sourceInternalAddress": <internalAddress>, "destinationLocation":
<location>, "destinationSN": < serialNumber >}]
```

Moves a cartridge to a slot chosen by the library:

```
POST /v1/workItems [{"type": "moveToSlot", "cartridge":
<volser>,"sourceInternalAddress": <internalAddress>}]
```

Moves a cartridge to an I/O station chosen by the library:

```
POST /v1/workItems [{"type": "moveToIOStation", "cartridge": <volser>,
"sourceInternalAddress": <internalAddress>}]
```

Move to drive

Add the ability to move cartridges to drives. To request movement of a cartridge, a work item is created in the work queue. The ID of the work item is returned after it is submitted:

```
POST /v1/workItems [{"type": "moveToDrive", "cartridge":
<volser>,"sourceInternalAddress": <internalAddress>, "destinationLocation":
<location>, "destinationSN": < serialNumber >}]
```

Request body parameters

The following parameters are available:

- `type`: A string that represents the type of work item that is being created (in this case, `moveToDrive`).
- `cartridge`: A string that represents the VOLSER of the cartridge to be moved. If duplicate VOLSERs exist in the library for this cartridge, the `sourceInternalAddress` is used instead.
- `sourceInternalAddress`: A string that represents the internal address of the location of the cartridge to be moved. This internal address might change if the cartridge is assigned or unassigned from a logical library or if the cartridge is moved by the host or system. This value must be queried immediately before use.
- `destinationLocation`: A string that represents the location of the slot (for example, "`slot_F1C3R23T0`"). Only tier 0 or tier 1 is allowed as a destination.
- `destinationInternalAddress`: A string that represents the internal address of the location to which the cartridge is to be moved.
- `cartridge`: A string that represents the VOLSER of the cartridge to be moved. If duplicate VOLSERs exist in the library for this cartridge, the `sourceInternalAddress` must be used instead.
- `sourceInternalAddress`: A string that represents the internal address of the location of the cartridge to be moved. This internal address might change if the cartridge is assigned or unassigned from a logical library or if the cartridge is moved by the host or system. This value must be queried immediately before use.
- `destinationLocation`: A string that represents the location of the slot (for example, "`slot_F1C3R23T0`"). Only tier 0 or tier 1 is allowed as a destination.
- `destinationInternalAddress`: A string that represents the internal address of the location to which the cartridge is to be moved.

Note: Some of the parameter settings include the following usage considerations:

- The source must be specified by entering the `cartridge` or `sourceInternalAddress`. Only one of these parameters is specified.
- The destination can be specified by the `destinationLocation`, the `destinationInternalAddress`, or neither. If neither is specified, the library chooses an open slot and an available internal address to move the cartridge to automatically.
- When moving a cartridge by using a `destinationLocation`, the command is rejected if the cartridge to be moved is in a tier 2 - 5 storage slot.
- When moving a data cartridge by using a `destinationInternalAddress`, the command is rejected if it is not moved to an internal address within the cartridge's logical library (and more complexity in the future). This parameter is not applicable to cleaning or diagnostic cartridges.

When the **`POST /v1/workItems {"type": "moveToDrive", "cartridge": <volser>, "destinationLocation": <location>`**`}` command runs successfully, the system returns output to your display, as shown in Example 6-36.

Example 6-36 POST command example

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg8<<<'POST /v1/workItems
{"type": "moveToDrive", "cartridge": "F20223L9", "destinationLocation":
"drive_F2C2R3"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg8
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 11
```

By using the TS4500 Web GUI to determine the status change of the drive in Frame 2, Column 2, Row 3, change the status to `ready` and cartridge with volser F20223L9 `loaded`. Ensure that the cartridge and the drive belong to the same logical library.

Move to slot

Add the ability to move cartridges to a slot that is chosen by the library. To request the movement of a cartridge, a work item is created in the work queue. The ID of the work item is returned after it is submitted:

```
POST /v1/workItems [{"type": "moveToSlot", "cartridge":
<volser>,"sourceInternalAddress": <internalAddress>}]
```

Request body parameters

The following parameters are available:

- `type`: A string that represents the type of work item that is being created (in this case, `moveToSlot`).
- `cartridge`: A string that represents the VOLSER of the cartridge to be moved. If duplicate VOLSERs exist in the library for this cartridge, the `sourceInternalAddress` must be used instead.

- `sourceInternalAddress`: A string that represents the internal address of the location of the cartridge to be moved. This internal address might change if the cartridge is assigned or unassigned from a logical library or if the cartridge is moved by the host or system. This value must be queried immediately before use.

> **Note:** The source must be specified by entering the `cartridge` or `sourceInternalAddress`. Only one of these parameters must be specified.

When the **`POST /v1/workItems {"type": "moveToSlot", "cartridge": <volser>}`** command runs successfully, the system returns output to your display, as shown in Example 6-37.

Example 6-37 Returned output

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg8<<<'POST /v1/workItems
{"type": "moveToSlot", "cartridge": "SGP396L3"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg8
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 11
```

If the cartridge that is being moved is in the I/O Station, a physical move into one of the storage slots occurs. The library chooses the destination. If a move for a cartridge that is in storage slot is to be done, only the element address is changed. No physical movement occurs.

Move to I/O Station

Add the ability to move cartridges to IO Station and the slot that is chosen by the library. To request movement of a cartridge, a work item is created in the work queue. The ID of the work item is returned after it is submitted:

```
POST /v1/workItems [{"type": "moveToIOStation", "cartridge": <volser>,
"sourceInternalAddress": <internalAddress>}]
```

Request body parameters

The following parameters are available:

- `type`: A string that represents the type of work item that is being created (in this case, `moveToIOStation`).
- `cartridge`: A string that represents the VOLSER of the cartridge to be moved. If duplicate VOLSERs exist in the library for this cartridge, the `sourceInternalAddress` must be used instead.
- `sourceInternalAddress`: A string that represents the internal address of the location of the cartridge to be moved. This internal address might change if the cartridge is assigned or unassigned from a logical library or if the cartridge is moved by the host or system. This value must be queried immediately before use.

> **Note:** The source must be specified by entering the `cartridge` or `sourceInternalAddress`. Only one of these parameters must be specified.

When the **`POST /v1/workItems {"type": "moveToIOStation", "cartridge": <volser>}`** command runs successfully, the system returns output to your display, as shown in Example 6-38.

Example 6-38 Returned output

```
/usr/bin/sg_write_buffer --id=0x10 --mode=2 --raw  /dev/sg8<<<'POST /v1/workItems
{"type": "moveToIOStation", "cartridge": "SGP396L3"}'

/usr/bin/sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --raw /dev/sg8
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 11
```

If a magazine does not exist in the I/O Station, or the magazine is full, the cartridge does not physically move but the internal address is changed to Virtual IO address. The status is changed to export pending.

6.3 Using ITDT for RoS

This section describes how to use IBM Tape Diagnostic Tool (ITDT) for RoS. We describe the formats and include examples by using the commands.

For more information about installing ITDT, see *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130.

Install ITDT-SE version 9.4 or later for your operating system.

Run the ITDT commands directly from the command line and include the path to the ITDT folder (if running from outside of the ITDT folder).

Command format

The ITDT commands feature the format that is shown in the following examples:

```
./itdt -f /dev/sgXX ros GET /v1/accessors/<location>
./itdt -f /dev/sgXX ros POST /v1/drives/<location>/clean
./itdt -f /dev/sgXX ros PATCH /v1/library/{"time":<time>}
```

Note: The **`ros`** parameter in the command string can be replaced with **`rosraw`** if raw (unformatted) data is required.

The Microsoft Windows command prompt does not support a single quotation mark (') to send a string to an application. Every special character must be escaped; for example:

```
ros POST /v1/tasks '{"type":"inventoryTier0and1","location":"library"}'
```

This command works on Linux, but must be sent on Windows as:

```
ros POST /v1/tasks {\"type\":\"inventoryTier0and1\",\"location\":\"library\"}
```

6.3.1 Examples

Example 6-39 shows the **ITDT** command with **ros** to get the Library information. The output is formatted.

Example 6-39 The ros GET /v1/library command

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros GET /v1/library
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 706
[
  {
    "name": "Anacg4",
    "status": "online",
    "totalCapacity": "1252",
    "licensedCapacity": "900",
    "totalCartridges": "637",
    "assignedCartridges": "71",
    "firmware": "1.6.0.1-B00",
    "sn": "78AA469",
    "time": "2019-11-20T03:48:12-0600",
    "location": "Reynholm Industries",
    "address": "123 Carenden Road",
    "city": "London",
    "state": "London",
    "country": "England",
    "contact": "Maurice Moss",
    "telephone": "0118-999-88199",
    "secondaryTelephone": "9119-725-3",
    "secureCommunications": "enabled",
    "autoEjectCleaningCartridges": "enabled",
    "elasticCapacity": "enabled",
    "activeAccessors": "Accessor A only",
    "vioStatus": "enabled",
    "nmaDetection": "enabled",
    "capacityUtilThresh": "99.0",
    "dualAccessorUtilThresh": "98.0"
  }
]

Exit with code: 0
```

Example 6-40 shows the **ITDT** command with **rosraw** to get the Library information. The output is unformatted.

Example 6-40 The rosraw GET /v1/library command

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 rosraw GET /v1/library
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 706
[{"name": "Anacg4", "status": "online", "totalCapacity": "1252",
"licensedCapacity": "900", "totalCartridges": "637", "assignedCartridges": "71",
"firmware": "1.6.0.1-B05", "sn": "78AA469", "time": "2019-11-20T03:49:54-0600",
"location": "Reynholm Industries", "address": "123 Carenden Road", "city":
```

```
"London", "state": "London", "country": "England", "contact": "Maurice Moss",
"telephone": "0118-999-88199", "secondaryTelephone": "9119-725-3",
"secureCommunications": "enabled", "autoEjectCleaningCartridges": "enabled",
"elasticCapacity": "enabled", "activeAccessors": "Accessor A only", "vioStatus":
"enabled", "nmaDetection": "enabled", "capacityUtilThresh": "99.0",
"dualAccessorUtilThresh": "98.0"}]??

Exit with code: 0
```


6.3.2 Endpoints

This section describes the endpoints and attributes that are defined in the REST API.

Accessors

Retrieve a list of all accessors or a single accessor resource:

```
GET /v1/accessors
GET /v1/accessors/<location>
```

Path parameters

The `location` path parameter is a string that describes the unique identifier for the accessor with a value of “`accessor_Aa`” or “`accessor_Ab`”.

When the **`GET /v1/accessors/accessor_Aa`** command runs successfully, the system returns output to your display, as shown in Example 6-41.

Example 6-41 GET /v1/accessor/accessor_Aa command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/accessors/accessor_Aa
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 231
[
  {
    "location": "accessor_Aa",
    "state": "onlineActive",
    "pivots": "3181",
    "barCodeScans": "20245",
    "travelX": "3988",
    "travelY": "3588",
    "getsGripper1": "3405",
    "putsGripper1": "3405",
    "getsGripper2": "3239",
    "putsGripper2": "3239"
  }
]

Exit with code: 0
```

The returned attribute values show the number of lifetime moves or metres that were travelled for each component. The “`state`” field returns the current state of the accessor. The possible states include the following values, which are listed in priority order:

- `inServiceMode`
- `noMovementAllowed`

- noMotorPower
- bothGrippersFailed
- gripper1Failed
- gripper2Failed
- calibrating
- onlineStandby
- onlineActive

Cleaning cartridges

Cleaning cartridges are used periodically to clean a drive. They have a limited number of uses.

The **GET /v1/cleaningCartridges** command retrieves a list of system managed cleaning cartridges or a single cleaning cartridge resource:

```
GET /v1/cleaningCartridges
GET /v1/cleaningCartridges/<volser>
```

Path parameters

The `volser` path parameter is a string that represents the volume serial number or bar code that uniquely identifies the cleaning cartridge; for example,"`CLN236JA`". If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

When the **GET /v1/cleaningCartridges** command runs successfully, the system returns output to your display, as shown in Example 6-42.

Example 6-42 Get /v1/cleaningCartridges command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/cleaningCartridges
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 1991
[{
    "volser": "CLN137JA",
    "state": "Normal",
    "cleansRemaining": "38",
    "location": "slot_F2C7R17T1",
    "mediaType": "3592",
    "mostRecentUsage": "2019-08-02T18:27:40+0000"
  },
{
    "volser": "CLN871JA",
    "state": "Normal",
    "cleansRemaining": "50",
    "location": "slot_F2C5R21T2",
    "mediaType": "3592",
    "mostRecentUsage": "2019-07-29T16:44:05+0000"
  }
]

Exit with code: 0
```

The following command attributes are available:

- `volser`: A string that represents the volume serial number or bar code that uniquely identifies the cartridge.

- `state`: A string that represents the current state of the cartridge. Values include:
  - `failedVerification`
  - `assignmentRequired`
  - `uncertainBarcode`
  - `exportQueued`
  - `importing`
  - `verifying`
  - `normal`
- `cleansRemaining`: A string that shows the number of remaining cleans.
- `location`: A string that describes the location of the cartridge.

  > **Note:** The gripper location is shown only as part of a failure scenario if the cartridge is stuck in the gripper or as part of a tier 2 mount. Otherwise, a cartridge that is in transit maintains its source location until after the move completes.

  The format depends on the location of the cartridge, as shown in the following examples:

  - gripper_A<“a”|“b”>G<”1”|”2”> (for example, “gripper_AbG1”)
  - drive_F<f>C<c>R<r> (for example, “drive_F1C3R23”)
  - slot_ F<f>C<c>R<r>T<t> (for example, “slot_F1C3R23T0”)
  - ioStationSlot_F<f>IO<”u”|”l”>S<s> (for example, “ioStationSlot_F1IOuS4”)
- `mediaType`: A string that represents the media type that is supported by the tape drive. Supported values include “3592” and “LTO”.
- `mostRecentUsage`: A string that represents the last date and time this cartridge was mounted into a drive, or null if this value is unknown or the cartridge is not yet mounted. Time format is in ISO 8601 format of “YYYY-MM-DDThh:mm:ss+-hmm”; for example, “2018-09-17T23:02:00-0700”.

> **Note:** Any cartridge with an “Unknown” or null VOLSER is assumed by the library to be a data cartridge; therefore, they never display in this query.

Data cartridges

Data cartridges retrieves information about all host-accessible or unassigned data cartridges in the tape library. If a data cartridge is removed from the library or missing during an inventory scan, it does not appear in this list:

```
GET /v1/dataCartridges
GET /v1/dataCartridges/<volser>
```

Path parameters

A string that represents the volume serial number or bar code that uniquely identifies the cartridge to a host; for example, “EZ1122L3”. If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

When the **`GET /v1/dataCartridges/`** command runs successfully, the system returns the output that is shown in Example 6-43.

Example 6-43 GET /v1/dataCartridges command

```
[root@TRON ITDT]# ./itdt -f /dev/sg4 ros GET /v1/dataCartridges
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 144210
```

```
},
  {
    "volser": "000002JC",
    "state": "normal",
    "accessible":"normal",
    "location": "slot_F2C5R22T1",
    "mediaType": "3592",
    "encrypted": null,
    "locked":"unlocked",
    "mostRecentVerification": null,
    "mostRecentUsage": "2020-01-30T09:54:52-0600",
    "logicalLibrary": null,
    "elementAddress": null
  },
  {
    "volser": "310517JA",
    "state": "normal",
    "accessible":"normal",
    "location": "slot_F2C10R36T0",
    "mediaType": "3592",
    "encrypted": null,
    "locked":"unlocked",
    "mostRecentVerification": null,
    "mostRecentUsage": "2020-01-30T10:02:31-0600",
    "logicalLibrary": null,
    "elementAddress": null
...
```

Note: In Example 6-43, the output should display *all* cartridges, but we show only two cartridges. When only the data is required, use the "`volser`" path parameter.

The following command attributes are available:

- `volser`: This string represents the volume serial number or bar code that uniquely identifies the cartridge to a host; for example, EZ0137JA. If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.
- `state`: A string that represents the current state of the cartridge. Values include the following in priority order:
  - `failed Verification`
  - `assignmentRequired`
  - `uncertainBarcode`
  - `exportQueued`
  - `importing`
  - `verifying`
  - `normal`
- `location`: A string that describes the location of the cartridge.

> **Note:** The gripper location is shown only as part of a failure scenario if the cartridge is stuck in the gripper or as part of a tier 2 mount. Otherwise, a cartridge that is in transit maintains its source location until after the move completes. The format depends on the location of the cartridge:
>
> - gripper_A<"a"|"b">G<"1"|"2"> (for example, "gripper_AbG1")
> - drive_F<f>C<c>R<r> (for example, "drive_F1C3R23")
> - slot_ F<f>C<c>R<r>T<t> (for example, "slot_F1C3R23T0")
> - ioStationSlot_F<f>IO<"u"|"l">S<s> (for example, "ioStationSlot_F1IOuS4")

- mediatype: A string that represents the media type that is supported by the tape drive. Supported values include "3592" and "LTO".
- encrypted: A string that represents the state of encryption on this cartridge. If "yes", the data on the cartridges was encrypted by the tape drive. If "no", the data on the cartridge was not encrypted by the tape drive. If null, the cartridge is not yet mounted.
- mostRecentVerification: A string that represents the last date and time this cartridge was verified by using media verification in a drive that is assigned as a media verification drive. If media verification was not run on this cartridge, this value returns as null.
- mostRecentUsage: A string that represents the last date and time this cartridge was mounted into a drive, or null if this value is unknown or the cartridge is not yet mounted. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm"; for example, "2020-03-17T13:02:00-0700".
- logicalLibrary: A string that represents the name of the logical library to which this cartridge is assigned. If it is not assigned to a logical library, this value is null. If the cartridge does not belong to a bar code range, it is available to all logical libraries. In this case, the state reported is "importing" and this value returns null.
- elementAddress: An integer that represents the element address of the cartridge drive. This value is reported to hosts by way of the **Read Element Status SCSI** command. This value is null if it is not yet reported to the host.

Diagnostic cartridges

Diagnostic cartridges are used to test tape drives.

The **GET /v1/diagnosticCartridges** command retrieves a list of diagnostic cartridges or a single diagnostic cartridge resource:

```
GET /v1/diagnosticCartridges
GET /v1/diagnosticCartridges/<volser>
```

Path parameters

The volser parameter is a string that represents the volume serial number or bar code that uniquely identifies the diagnostic cartridge; for example, "DG 031L5". If duplicate VOLSERs exist in the library, a query of this VOLSER returns more than one cartridge.

When the **GET /v1/diagnosticCartridges** command runs successfully, the system returns output to your display, as shown in Example 6-44.

Example 6-44 GET /v1/diagnosticCartridges command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/diagnosticCartridges
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 1413
```

```
[
  {

    "volser": "CE 002JL",
    "state": "Normal",
    "location": "slot_F1C4R3T0",
    "mediaType": "3592",
    "mostRecentUsage": "2019-11-06T19:35:37+0000"
  },
{
    "volser": "CE 882JA",
    "state": "Normal",
    "location": "ioStationSlot_F2OuS1",
    "mediaType": "3592",
    "mostRecentUsage": "2019-11-11T14:02:19+0000"
  }
]
Exit with code: 0
```

The attributes are the same as those attributes for cleaning cartridges.

Note: Any cartridge with an "Unknown" or null VOLSER is assumed by the library to be a data cartridge; therefore, they never show up in this query.

Get drives

Retrieve a list of drives or a single drive resource.

The **`GET /v1/drives`** command is used to obtain drive information for all drives or a single drive resource:

```
GET /v1/drives
GET /v1/drives/<location>
GET /v1/drives/<sn>
```

Parameters

The following parameters are available:

- `location`: A string that represents the unique location of the drive. This value also is the unique identifier for the drive. The format is *drive_F<f>C<c>R<r>">"* where **"**f**"** is the frame, **"**c**"** is the column, and **"**r**"** is the row in which the drive is installed.
- `sn`: A string that represents the serial number of the drive. This value can also be used as a unique identifier for the tape drive.

When the **`GET /v1/drives`** command runs successfully, the system returns output to your display, as shown in Example 6-45.

Example 6-45 GET /v1/drives command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/drives/drive_F2C4R4
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 417
[
  {
    "location": "drive_F2C4R4",
```

```
    "sn": "0000078PG204",
    "mediaType": "3592",
    "state": "online",
    "operation": "empty",
    "mtm": "3592-60F",
    "Bar code": "11S02CE664YF1013000638",
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "JAG6_2",
    "use": "controlPath",
    "firmware": "559B",
    "encryption": "disabled",
    "wwnn": "500507604419ff0f",
    "elementAddress": "264",
    "beacon": "disabled",
    "volser": null
  }
]
Exit with code: 0
```

The following command attributes are available:

- `mediaType`: A string that represents the media type teeth is supported by the tape drive. Supported values include **"**`3592`**"** and **"**`LTO`**"**.
- `state`: A string that represents the current state of the tape drive. Values include the following values in priority order:
  - `inServiceMode`
  - `restarting`
  - `initializing`
  - `unreachable`
  - `resetRequired`
  - `updating`
  - `cleaning`
  - `online`
- `operation`: A string that represents the current operation the tape drive is performing. If the tape drive is not in the "online" state, this value is null. Values include the following values in priority order:
  - `empty`: Drive is online and empty.
  - `loading`: Drive is online and in the process of loading a cartridge.
  - `unloading`: Drive is online and unloading writing to the loaded cartridge.
  - `reading`: Drive is online and reading from the loaded cartridge.
  - `writing`: Drive is online and writing to the loaded cartridge.
  - `locating`: Drive is online and doing a seek action on the loaded cartridge.
  - `rewinding`: Drive is online and rewinding the loaded cartridge.
  - `erasing`: Drive is online and erasing the loaded cartridge.
  - `formatting`: Drive is online and formatting the loaded cartridge.
  - `calibrating`: Drive is online and calibrating.
  - `ready`: Drive is online, shows no activity, and has a cartridge mounted. This value many times matches the drive's SCSI ready condition, but not always.

- unloaded: Drive is online, shows no activity, and has no cartridge mounted.

- mtm: A string that represents the machine type and model of the tape drive (for example, 3588-F6C).
- Bar code: A string that represents the 11S bar code label found on the tape drive. This value is a 22-character string.
- interface: A string that represents the type of ports this tape drive includes. Values include:
  - fibreChannel
  - iSCSI
- interfaceMode: A string that represents the interface mode supported by the ports this tape drive includes. Values include: "single-mode" for single-mode fiber and "multi-mode" for multi-mode fiber. The value is null if the ports included have an iSCSI interface.
- logicalLibrary: A string that represents the name of the logical library this tape drive is assigned to or null if the tape drive is unassigned.
- use: A string that represents the assigned use this tape drive was given within the logical library. Values include:
  - access: Data access drives
  - controlPath: Data access and control path drives
  - verification: Media verification drives
- firmware: A string that represents the firmware version that is installed on the tape drive.
- encryption: The state of encryption on this tape drive. Values include:
  - enabled
  - disabled
- wwnn: A 16-character hex string that represents the worldwide node name of the tape drive.
- elementAddress: An integer that represents the element address of the tape drive. This value is reported to hosts by way of the **Read Element Status SCSI** command. This value is null if unassigned to a logical library.
- beacon: A string indicating the current state of the drive beacon LED. This beacon can be seen from the rear of the drive. Values include:
  - enabled (LED is flashing)
  - disabled (LED is off)
- volser: A string that represents the volume serial number or bar code that uniquely identifies the cartridge that is mounted in the tape drive; for example, "SG1122L2" or null if the tape drive does not have a cartridge mounted or the loaded cartridge's VOLSER is unknown.

Note: If the drive is replaced while the beacon is enabled, the state of the physical LED might not match this attribute. In this case, the beacon should be disabled.

Clean drives

Run a drive clean on a specific drive:

```
POST /v1/drives/<location>/clean
```

Path parameters

The `location` parameter is a string that represents the location of the tape drive. This value also is the unique identifier for the drive. The format is **"**`drive_F<f>C<c>R<r>`**"** where **"**`f`**"** is the frame, **"**`c`**"** is the column, and **"**`r`**"** is the row in which the drive is installed.

When the **`POST /v1/drives`** command runs successfully, the system returns output to your display, as shown in Example 6-46.

Example 6-46 POST /v1/drives/drive_F1C2R3/clean command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros POST /v1/drives/drive_F1C2R3/clean
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 0
Exit with code: 0
```

The command prompt returns without comments. However, if we check the status of `drive_F1C3R3` before and after command execution, we see that a cleaning cartridge was mounted and the drive is performing a cleaning cycle.

Before command execution is shown in Example 6-47.

Example 6-47 GET /v1/drives/drive_F1C2R3 command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/drives/drive_F1C2R3
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 395
[
  {
    "location": "drive_F1C2R3",
    "sn": "0000013B004C",
    "mediaType": "3592",
    "state": "online",
    "operation": "empty",
    "mtm": "3592-EH7",
    "Bar code": null,
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "JAG4",
    "use": "access",
    "firmware": "3DCB",
    "encryption": "disabled",
    "wwnn": "500507604419fff6",
    "elementAddress": "259",
    "beacon": "disabled",
    "volser": null
  }
]

Exit with code: 0
```

After command execution is shown in Example 6-48.

Example 6-48 GET /v1/drives/drive_F1C2R3 command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/drives/drive_F1C2R3
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 400
[
  {
    "location": "drive_F1C2R3",
    "sn": "0000013B004C",
    "mediaType": "3592",
    "state": "cleaning",
    "operation": null,
    "mtm": "3592-EH7",
    "Bar code": null,
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "JAG4",
    "use": "controlPath",
    "firmware": "3DCB",
    "encryption": "disabled",
    "wwnn": "500507604419fff6",
    "elementAddress": "259",
    "beacon": "disabled",
    "volser": "CLN481JA"
  }
]
Exit with code: 0
```

Modify drive use

Modify the use of the drive to be a data access, control path, or verification drive:

```
PATCH /v1/drives/<location> {"use": <"access" | "controlPath" | "verification">}
```

Path parameters

The `location` parameter is a string that represents the location of the tape drive. This value also is the unique identifier for the drive. The format is "`drive_F<f>C<c>R<r>`" where "`f`" is the frame, "`c`" is the column, and "`r`" is the row in which the drive is installed.

Request body parameters

The `use` parameter is a string that represents the assigned use this tape drive was given within the logical library. Values include:

- `access`: data access drive
- `controlPath`: data access and control path drive
- `verification`: media verification drive

When the **`PATCH /v1/drives`** command runs successfully, the system returns output to your display, as shown in Example 6-49.

Example 6-49 PATCH /v1/drives/drive_F1C2R3 {"use": "controlPath"} command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros PATCH /v1/drives/drive_F1C2R3 {"use":
"controlPath"}
HTTP/1.1 200 OK
Content-Type: application/json
```

```
Content-Length: 0
Exit with code: 0
```

The command prompt returns without comments. However, if we check the status of the `drive_F2C3R1` before and after command execution, we see that it was changed to controlPath usage.

Before command execution is shown in Example 6-50.

Example 6-50 GET /v1/drives/drive_F1C2R3 command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/drives/drive_F1C2R3
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 395
[
  {
    "location": "drive_F1C2R3",
    "sn": "0000013B004C",
    "mediaType": "3592",
    "state": "online",
    "operation": "empty",
    "mtm": "3592-EH7",
    "Bar code": null,
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "JAG4",
    "use": "access",
    "firmware": "3DCB",
    "encryption": "disabled",
    "wwnn": "500507604419fff6",
    "elementAddress": "259",
    "beacon": "disabled",
    "volser": null
  }
]

Exit with code: 0
```

After command execution is shown in Example 6-51.

Example 6-51 GET /v1/drives/drive_F1C2R3 command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/drives/drive_F1C2R3
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 395
[
  {
    "location": "drive_F1C2R3",
    "sn": "0000013B004C",
    "mediaType": "3592",
    "state": "online",
    "operation": "empty",
    "mtm": "3592-EH7",
```

```
    "Bar code": null,
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "JAG4",
    "use": "controlPath",
    "firmware": "3DCB",
    "encryption": "disabled",
    "wwnn": "500507604419fff6",
    "elementAddress": "259",
    "beacon": "disabled",
    "volser": null
  }
]

Exit with code: 0
```

Turn drive beacon on/off

Turns on and off the drive beacon LED at the back of the drive:

`PATCH /v1/drives/<location> {“beacon”: <”enabled” | “disabled”>}`

Path parameters

The `location` parameter is a string that represents the location of the tape drive. This value 0 is also the unique identifier for the drive. The format is “`drive_F<f>C<c>R<r>`” where “`f`” is the frame, “`c`” is the column, and “`r`” is the row in which the drive is installed.

Request body parameters

The `beacon` parameter is a string that indicates the current state of the drive beacon LED. This beacon can be seen from the rear of the drive. Values include: “`enabled`” (LED is flashing) and “`disabled`” (LED is off).

Note: If the drive is replaced while the beacon is enabled, the state of the physical LED might not match this attribute. In this case, the beacon should be disabled.

When the **`PATCH /v1/drives/drive_F1C2R3 {"beacon": "enabled"}`** command runs successfully, the system returns output to your display, as shown in Example 6-52.

Example 6-52 PATCH /v1/drives/drive_F1C2R3 {"beacon": "enabled"} command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros PATCH /v1/drives/drive_F1C2R3 {"beacon":
"enabled"}
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 0
Exit with code: 0
```

The command prompt returns without comments. However, if we check the status of the drive_F1C2R3 before and after the command is run, we see that the beacon changed from “`disabled`” to “`enabled`”:

Before command execution is shown in Example 6-53.

Example 6-53 GET /v1/drives/drive_F1C2R3 command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/drives/drive_F1C2R3
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 395
[
  {
    "location": "drive_F1C2R3",
    "sn": "0000013B004C",
    "mediaType": "3592",
    "state": "online",
    "operation": "empty",
    "mtm": "3592-EH7",
    "Bar code": null,
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "JAG4",
    "use": "controlPath",
    "firmware": "3DCB",
    "encryption": "disabled",
    "wwnn": "500507604419fff6",
    "elementAddress": "259",
    "beacon": "disabled",
    "volser": null
  }
]

Exit with code: 0
```

After command execution is shown in Example 6-54.

Example 6-54 GET /v1/drives/drive_F1C2R3 command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/drives/drive_F1C2R3
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 395
[
  {
    "location": "drive_F1C2R3",
    "sn": "0000013B004C",
    "mediaType": "3592",
    "state": "online",
    "operation": "empty",
    "mtm": "3592-EH7",
    "Bar code": null,
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "JAG4",
    "use": "controlPath",
    "firmware": "3DCB",
    "encryption": "disabled",
    "wwnn": "500507604419fff6",
```

```
    "elementAddress": "259",
    "beacon": "enabled",
    "volser": null
  }
]
Exit with code: 0
```

Get Ethernet ports

Retrieve a list of Ethernet ports or a single Ethernet port resource:

```
GET /v1/ethernetPorts
GET /v1/ethernetPorts/<location>
```

Path parameters

The `location` parameter is a string that represents the unique location of this port. The format of this field is "`eithernetPort_F<f>P<p>`" with possible values for `<p>` of "`a`", "`b`", "`imc`", "`tssc`", or "`service`".

When the **`GET /v1/ethernetPorts/<location>`** command runs successfully, the system returns output to your display, as shown in Example 6-55.

Example 6-55 GET /v1/ethernetPorts/ethernetPort_F1Pa command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/ethernetPorts/ethernetPort_F1Pa
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 497
[
  {
    "location": "ethernetPort_F1Pa",
    "macAddress": "40:F2:E9:52:71:85",
    "ipv4Address": "9.18.77.45",
    "ipv4Subnet": "255.255.255.0",
    "ipv4Gateway": "9.18.77.1",
    "ipv4Assignment": "static",
    "ipv4Primary": "9.0.128.50",
    "ipv4Secondary": "9.0.130.50",
    "ipv6Address": "disabled",
    "ipv6PrefixLength": "disabled",
    "ipv6Gateway": "disabled",
    "ipv6Primary": "disabled",
    "ipv6Secondary": "disabled",
    "ipv6Link": "disabled",
    "ipv6DHCP": "disabled",
    "ipv6StatelessConfig": "disabled",
    "ipv6Static": "disabled"
  }
]

Exit with code: 0
```

The following command attributes are available:

- `location`: A string that represents the unique location of this port. The format of this field is "`ethernetPort_F<f>P<p>`" with possible values for `<p>` of "`a`", "`b`", "`imc`", "`tssc`", or "`service`".
- `macAddress`: The media access control (MAC) address of this Ethernet port.
- `ipv4Address`: The IPv4 address of this Ethernet port.
- `ipv4Subnet`: The IPv4 subnet mask of this Ethernet port.
- `ipv4Gateway`: The IPv4 gateway address of this Ethernet port.
- `ipv4Assignment`: The IPv4 address of this Ethernet port; values include:
  - `static`
  - `dynamic`
- `ipv4Primary`: The IPv4 primary DNS address of this Ethernet port. If the ipv4Assignment is "`static`", this value is null.
- `ipv4Secondary`: The IPv4 secondary DNS address of this Ethernet port. If the ipv4Assignment is "`static`", this value is null.
- `Ipv6Address`: The IPv6 address of this Ethernet port.
- `ipv6PrefixLength`: The IPv6 prefix length of this Ethernet port.
- `ipv6Gateway`: The IPv6 gateway address of this Ethernet port.
- `ipv6Primary`: The IPv6 primary DNS address of this Ethernet port.
- `ipv6Secondary`: The IPv6 secondary DNS address of this Ethernet port.
- `ipv6Link`: The IPv6 link address of this Ethernet port.
- `ipv6DHCP`: The IPv6 DHCP setting of this Ethernet port.
- `ipv6StatelessConfig`: The IPv6 stateless config setting of this Ethernet port.
- `ipv6Static`: The IPv6 static address of this Ethernet port.

Get events

Retrieve a list of events or a single event resource. This command includes advanced filtering because many events can be in a library:

```
GET /v1/events
GET /v1/events/<ID>
GET /v1/events?{<"after=time" | "before=time" | "location">}
```

Path parameters

The `ID` parameter is a decimal integer that is the unique identifier of the event in the library.

Query parameters

The following parameters are available:

- `after (Optional)`: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`"; for example, "`2018-09-17T23:02:00`". If the time zone is omitted, the system's current time zone is used.

- before (Optional): A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of "YYYY-MM-DDThh:mm:ss+-hhmm" or "YYYY-MM-DDThh:mm:ss"; for example, "2018-09-17T23:02:00". If the time zone is omitted, the system's current time zone is used.
- location (Optional): A string that represents the hardware component where the event originated. Hardware components include:
  - accessor_A<"a"|"b"> (for example, "accessor_Ab")
  - gripper_A<"a"|"b">G<"1"|"2"> (for example, "gripper_AbG1")
  - column_F<f>C<c> (for example, "column_F1C3")
  - frame_F<f> (for example, "frame_F1")
  - fiducial_F<f>C<c>L<"t"|"b"> (for example, "fiducial_F1C3IOuLt")
  - fiducial_F<f>IO<"u"|"l">L<"t"|"b"> (for example, "fiducial_F1IOuLb")
  - drive_F<f>C<c>R<r> (for example, "drive_F1C3R23")
  - slot_ F<f>C<c>R<r>T<t> (for example, "slot_F1C3R23T0")
  - ioStation_F<f>IO<"u"|"l"> (for example, "ioStation_F1IOu")
  - ethernetPort_F<f>P<p> (for example, "ethernetPort_F2Pa").
  - frameSide_F<f>S<"A"|"B"> (for example, "frameSide_F1Sa")
  - position_F<f>C<c>R<r> (for example, "position_F1C5R7")

When the **GET /v1/events?<after>** command runs successfully, the system returns output to your display, as shown in Example 6-56.

Example 6-56 GET /v1/events?after=2019-11-12T06:22:15 command

```
[root@TRON ITDT]# [root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET
/v1/events?after=2019-11-12T06:22:15
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 1398
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 936
[
  {
    "ID": "8337",
    "severity": "information",
    "time": "2019-11-12T10:40:50-0600",
    "type": "audit",
    "location": "position_F1C2R3",
    "user": "SCSIControlPath_F2C2R4",
    "description": "A drive beacon was turned on",
    "state": null,
    "errorCode": "0820"
  },
  {
    "ID": "8336",
    "severity": "information",
    "time": "2019-11-12T10:08:40-0600",
    "type": "audit",
    "location": null,
    "user": "admin",
    "description": "A user logged in to the GUI from 9.145.82.93",
    "state": null,
    "errorCode": "0800"
```

```
  },
  {
    "ID": "8335",
    "severity": "information",
    "time": "2019-11-12T10:07:08-0600",
    "type": "audit",
    "location": null,
    "user": "admin",
    "description": "A user logged out of the GUI from 9.145.82.93",
    "state": null,
    "errorCode": "0812"
  },
  {
    "ID": "8334",
    "severity": "information",
    "time": "2019-11-12T10:07:07-0600",
    "type": "audit",
    "location": null,
    "user": "admin",
    "description": "A user logged out of the GUI from 9.86.127.31",
    "state": null,
    "errorCode": "0812"
  }
]

Exit with code: 0
```

The following command attributes are available:

- `ID`: A decimal integer that is the unique identifier of the event in the library.
- `severity`: A string that represents the severity of the event that indicates the urgency that should be given to it. Values include:
  - `error`
  - `warning`
  - `inactiveError`
  - `inactiveWarning`
  - `information`
- `time`: The time that the event occurred. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`”; for example, “`2018-09-17T23:02:00-0700`”.
- `location`: A string that represents the hardware component where the event originated. This value can be null if no hardware was associated with this event. Hardware components include:
  - `accessor_A<“a”|“b”>`
  - `gripper_A<“a”|“b”>G<”1”|”2”>`
  - `frame_F<f>`
  - `fiducial_F<f>C<c>L<”t”|”b”>`
  - `fiducial_F<f>IO<”u”|”l”>L<”t”|”b”>`
  - `drive_F<f>C<c>R<r>`
  - `slot_ F<f>C<c>R<r>T<t>`
  - `ioStation_F<f>IO<”u”|”l”> )`
  - `ethernetPort_F<f>P<p>`
  - `frameSide_F<f>S<”A”|”B”>`
  - `position_F<f>C<c>R<r>`

- `user`: A string that shows the user name of the user that issued the action that caused this event. This value can be null if this event did not occur as a result of a user action.
- `description`: A string that describes the event that gives more information about exactly what it represents.
- `state`: A string that represents the current state of the event regarding its Call Home status. This value can be null if no Call Home occurred yet. Values include:
  - `Detected error <error code>`"
  - `Command failed with error code <error code>`"
  - `Calling home`"
  - `Call home initiated`"
  - `Call home failed`"
  - `Call home not attempted because not configured`"
  - `Assigned PMR <PMR number>. Service action required`"
  - `Service action complete by <description>`"
- `errorCode`: A 4-digit hex string that represents the internal error code that is used to represent this type of event.

Get Frames

General information of the installed frames of the library.

Retrieve a list of the currently installed frames of the library:

```
GET /v1/frames
GET /v1/frames/<location>
```

Path parameters

The `location` parameter is a string that represents the unique location of the frame in the library. The format is "`frame_F<f>`" where `'f'` is the frame position.

When the **`GET /v1/frames`** command runs successfully, the system returns output to your display, as shown in Example 6-57.

Example 6-57 Get /v1/frames command

```
[root@TRON ITDT]# ./itdt -f /dev/sg4 ros GET /v1/frames
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 405
[
  {
    "location": "frame_F1",
    "state": "calibrationRequired",
    "type": "Expansion frame",
    "mtm": "3584-D25",
    "sn": "78CB254",
    "mediaType": "3592",
    "frontDoor": "closed",
    "frontDoorLastChanged": "2021-10-18T14:18:33+0000",
    "rearDoor": null,
    "rearDoorLastChanged": null,
    "sideDoor": "closed",
    "sideDoorLastChanged": "2021-05-21T20:24:45+0000",
    "slots": 590,
    "cartridges": 354,
    "drives": 9,
```

```
    "ioStations": 0
  },
  {
    "location": "frame_F2",
    "state": "calibrationRequired",
    "type": "Base frame",
    "mtm": "3584-L25",
    "sn": "78AB117",
    "mediaType": "3592",
    "frontDoor": "closed",
    "frontDoorLastChanged": "2021-10-18T14:18:33+0000",
    "rearDoor": null,
    "rearDoorLastChanged": null,
    "sideDoor": "closed",
    "sideDoorLastChanged": "2021-05-21T20:24:45+0000",

    "slots": 660,
    "cartridges": 265,
    "drives": 8,
    "ioStations": 2
  }
]

Exit with code: 0
```

The following command attributes are available:

- location: A string that represents the unique location of the frame in the library. The format is “frame_F<f>” where 'f' is the frame position.
- state: A string that represents the current state of the frame. Values include the following in priority order:
  - frontDoorOpenWhileNotAllowed
  - frontDoorOpen
  - sideDoorOpen
  - acUnreachable
  - calibrationRequired
  - inventoryPending
  - normal
- type: A string that represents the type of the frame.
- mtm: A string that represents the machine type and model of the frame. Values include:
  - Base frame
  - Expansion frame
  - Storage-only expansion frame
- mtm: A string that represents the machine type and model of this frame.
- sn: A string that represents the unique serial number of the frame.
- mediaType: A string that represents the media type that is supported by the tape drive. Supported values include “3592” and “LTO”.
- frontDoorString (open, closed)
- frontDoorLastChangedString (time/date)
- rearDoorString (open, closed)

- `rearDoorLastChangedString` (time/date)
- `sideDoorString` (open, closed)
- `sideDoorLastChangedString` (time/date)
- `slots`: A string that represents the number of available physical slots in this frame.
- `cartridges:` A string that represents the number of cartridges that are in slots or drives in this frame.
- `drives:` A string that represents the number of LTO drives that are installed in this frame.
- `ioStations`: Represents the number of I/O stations that are installed in this frame.

Get library

Retrieve the library resource.

The **`GET /v1/library`** command is used to obtain general information about the TS4500 tape library and its settings:

```
GET /v1/library
```

When the **`GET /v1/library`** command runs successfully, the system returns output to your display, as shown in Example 6-58.

Example 6-58 GET /v1/library command

```
[root@TRON ITDT]#  ./itdt -f /dev/sg34 ros GET /v1/library
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 706
[
  {
    "name": "Anacg4",
    "status": "online",
    "totalCapacity": "1252",
    "licensedCapacity": "900",
    "totalCartridges": "637",
    "assignedCartridges": "72",
    "firmware": "1.6.0.1-B00",
    "sn": "78AA469",
    "time": "2019-11-12T11:34:36-0600",
    "location": "Reynholm Industries",
    "address": "123 Carenden Road",
    "city": "London",
    "state": "London",
    "country": "England",
    "contact": "Maurice Moss",
    "telephone": "0118-999-88199",
    "secondaryTelephone": "9119-725-3",
    "secureCommunications": "enabled",
    "autoEjectCleaningCartridges": "enabled",
    "elasticCapacity": "enabled",
    "activeAccessors": "Accessor A only",
    "vioStatus": "enabled",
    "nmaDetection": "enabled",
    "capacityUtilThresh": "99.0",
    "dualAccessorUtilThresh": "98.0"
  }
```

```
]

Exit with code: 0
```

The following command attributes are available:

- `name`: A string that represents the name the tape library was given.
- `status`: A string that represents the overall health status of the library. Values include:
  - `doorOpenWhileNotAllowed`
  - `doorOpen`
  - `pausing`
  - `paused`
  - `restarting`
  - `inServiceMode`
  - `accessorsUnavailable`
  - `accessorDegraded`
  - `nodeCardDegraded`
  - `updating`
  - `scanningInventory`
  - `online`
- `totalCapacity`: An integer that represents the total physical cartridge capacity of the library.
- `licensedCapacity`: An integer that represents the total licensed cartridge capacity of the library.
- `totalCartridges`: An integer that represents the total number of cartridges in the library slots, I/O stations, drives, and accessors currently.
- `assignedCartridges`: An integer that represents the total number of cartridges assigned to logical libraries in the library.
- `firmware`: A string that represents the firmware level that is installed on the library.
- `sn`: A string that represents the serial number of the library.
- `time`: A string that represents the current date and time set on the library. All dates and times returned by the library are in relation to this time. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`". For example, "`2018-09-17T23:02:00-0700`".
- `location`: A string showing the user entered physical location of the library. For example, "`Building 9062`". If not provided, this is null.
- `address`: A string that represents the physical address of the library. If not provided, this is null.
- `city`: A string that represents the city in which the library is located. If not provided, this is null.
- `state`: A string that represents the two-character state in which the library is located. For example, "`AZ`" for Arizona. If not provided, this is null.
- `country`: A string that represents the country in which the library is located. If not provided, this is null.
- `contact`: A string that represents the name of the primary contact for the tape library. This string is used by IBM Support during repair and service. If not provided, this value is null.
- `telephone`: A string that represents the primary telephone number for the primary contact. This string is used by IBM Support during repair and service. If not provided, this value is null.

- `secondaryTelephone`: A string that represents the secondary telephone number for the primary contact. This string is used by IBM Support during repair and service. If not provided, this value is null.
- `secureCommunications`: A string that represents the current setting for secure communications. If “`enabled`”, the user is required to log in to the GUI by using a secure HTTPS connection. If “disabled”, an unsecured HTTP connection is allowed.
- `autoEjectCleaningCartridges`: A string that represents the current setting for automatically ejecting cleaning cartridges to the I/O Station. If “`enabled`”, cleaning cartridges are auto-ejected from the library when they have no cleans remaining. If “`disabled`”, they must be ejected manually.
- `elasticCapacity`: A string that represents the current setting for elastic capacity:
  - `max Capacity`
  - `temp Overflow`
  - `donates`
- `activeAccessors`: A string that represents which accessors are currently active. Values include:
  - `dualActive`
  - `accessorAOnly`
  - `accessorBOnly`
- `vioStatus`: A string that represents the current setting for Virtual I/O (VIO). This is “`enabled`” or “`disabled`”.
- `nmaDetection`: A string that represents the current setting for NMA detection. This is “`enabled`” or “`disabled`”.
- `capacityUtilThresh`: A percentage value accurate to one decimal place.
- `dualAccessorUtilThresh`: A percentage value accurate to one decimal place.

Reset library

Run a library reset.

The **`POST /v1/library/reset`** command is used to reset the entire library.

```
POST /v1/library/reset
```

When the **`POST /v1/library/reset`** runs successfully the system returns output to your display, as shown in Example 6-59.

Example 6-59 POST /v1/library/reset command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros POST /v1/library/reset
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 0
Exit with code: 0
```

Save Configuration (Data Base Backup)

The **`GET /v1/library/saveConfig`** command saves the library configuration to the same directory from which you run the command. This configuration file can be used to restore the machine to a previous state in case of need.

When the **GET /v1/library/saveConfig** command is issued, the result is as shown in Example 6-60.

Example 6-60 GET/v1/library/saveConfig

```
/./itdt -f /dev/sg8 ros GET /v1/library/saveConfig

File exported to: TS4500_SAVECONF_20211026010740.dbz

[root@tron ITDT]# ls -rtl *dbz
-rw-r--r-- 1 root root 4760529 Oct 26 01:08 TS4500_SAVECONF_20211026010740.dbz
```

Set library time, date, and time zone

Set the library time, date, and time zone:

```
PATCH /v1/library {“time”: <time>}
```

Request body parameter

The `time` parameter is the current date and time to be set on the library. All dates and times that are returned by the library are in relation to this time. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`”; for example, “`2018-09-17T23:02:00-0700`”.

When the **PATCH /v1/library {“time”: <time>}** command runs successfully, the system returns output to your display, as shown in Example 6-61.

Example 6-61 PATCH /v1/library {"time": "2019-11-12T10:09:39-0600"} command

```
./itdt -f /dev/sg34 ros PATCH /v1/library {"time": "2019-11-12T10:09:39-0600"}
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 0
Exit with code: 0
```

You can confirm the change by running the **GET /v1/library** command again. Alternatively, you can check the GUI by selecting **Settings** → **Library** → **Date and Time**. The date and time is the same as set by the ITDT command with a time zone of -06:00h (see Figure 6-2).

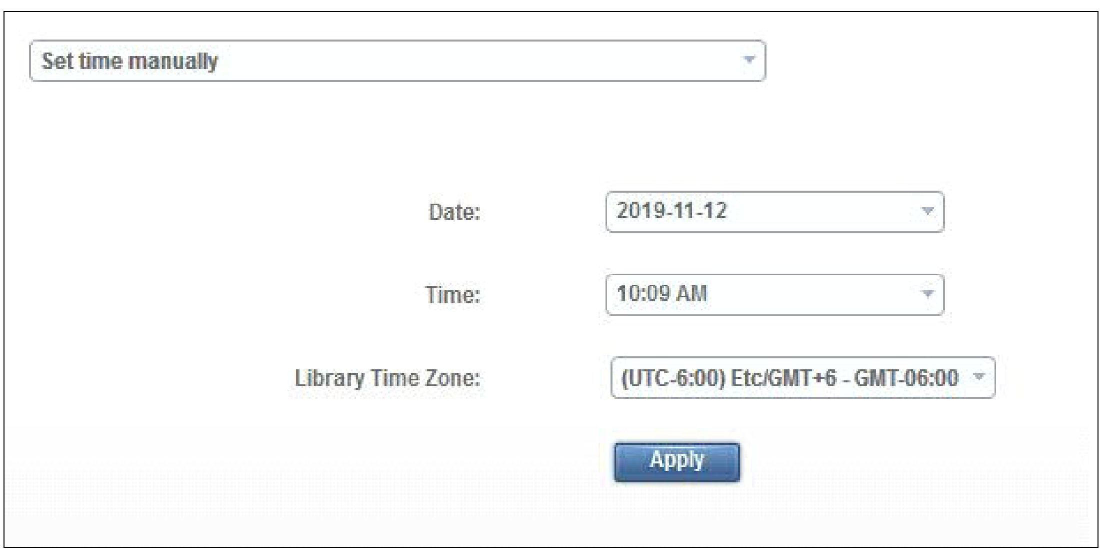

Figure 6-2 GUI - Library Date and Time

Logs

Retrieve a list of library logs, create a library log or export a created library log:

```
GET /v1/logs
POST /v1/logs
GET /v1/logs/<filename>
GET /v1/logs/<filename>/export
```

Path parameter

The `filename` parameter is the name of the file that is created or in the process of being created. This file name contains the name of the library, the date and time it was taken, and other important service information.

The **`GET /v1/logs`** command displays the list of previously created library logs. The output is similar to the output that is shown in Example 6-62.

Example 6-62 GET /v1/logs command output

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/logs
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 752
[
{
    "filename": "TS4500_LOG_AA469_201107093211.zip",
    "state": "completed",
    "startTime": "2019-11-07T09:32:11-0600",
    "lastUpdateTime": "2019-11-07T09:43:36-0600",
    "percentComplete": "100"
  },
  {
    "filename": "TS4500_LOG_AA469_20191105104707.zip",
    "state": "completed",
    "startTime": "2019-11-05T10:47:07-0600",
    "lastUpdateTime": "2019-11-05T10:55:23-0600",
    "percentComplete": "100"
  }
]
Exit with code: 0
```

The **`POST /v1/logs`** command is used to create a library log for IBM Support. When the **`POST /v1/logs`** command runs successfully the system returns output to your display, as shown in Example 6-63.

Example 6-63 POST /v1/logs command output

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros POST /v1/logs
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 52
{
  "filename": "TS4500_LOG_AA469_20191126074001.zip"
}Exit with code: 0
```

The library log that is created by running the **POST /v1/logs** command can be monitored by running the **GET /v1/logs<filename>** command. We can see **"state": "inProgress",** as shown in Example 6-64.

Example 6-64 GET /v1/logs<filename> command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET
/v1/logs/TS4500_LOG_AA469_20191126074001.zip
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 186
[
  {
    "filename": "TS4500_LOG_AA469_20191126074001.zip",
    "state": "inProgress",
    "startTime": "2019-11-26T07:40:01-0600",
    "lastUpdateTime": "2019-11-26T07:41:12-0600",
    "percentComplete": 40
  }
]

Exit with code: 0
```

The following command attributes are available:

- `state`: A string that shows the completion status of the event. Values include:
  - `inProgress`: The task is in progress.
  - `completed`: The task completed successfully.
  - `failed`: The task failed. Retry the attempt. If it continues to fail, contact IBM Support.

The **GET /v1/logs/<filename>/export** command can be used to export the library to the directory you run the command.

The use of the **GET /v1/logs/<filename>/export** is shown in Example 6-65.

Example 6-65 GET /v1/logs/filename/export command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET
/v1/logs/TS4500_LOG_AA469_20191105104707.zip/export
HTTP/1.1 200 OK
Content-Type: application/zip
Content-Length: 119808705

PK
Exit with code: 0
```

Get node cards

Retrieve a list of node cards or a single node card resource:

```
GET /v1/nodeCards
GET /v1/nodeCards/<ID>
POST /v1/nodeCards/<ID>/reset
```

Path parameter

The `ID` parameter is a decimal integer that is the unique identifier of the node card in the library.

When the **GET /v1/nodeCards/<ID>** command runs successfully, the system returns output to your display as shown in Example 6-66.

Example 6-66 GET /v1/nodeCards/17 command output

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/nodeCards/17
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 285
[
  {
    "ID": "17",
    "type": "MDA",
    "location": "accessor_Aa",
    "state": "online",
    "partNum": "38L7590",
    "sn": "Y010MY95Z61X",
    "Bar code": "11S38L7590Y010MY95Z61X",
    "ec": "0",
    "firmware": "0068",
    "cfBarcode": null,
    "cfPartNum": null,
    "cfVendor": null,
    "reportingLCC": null,
    "primaryLCC": null
  }
]

Exit with code: 0
```

The following command attributes are available:

- `ID`: A decimal integer that is the unique identifier of the node card in the library.
- `type`: A string that represents the physical type of the node card. Values include:
  - `MDA`: Motor driver assembly
  - `ACC`: Accessor control card
  - `LCC`: Library control card
- `location`: A string that represents the hardware component where this node card is physically located. For the MDA card and ACC, this string identifies the accessor in the format "`accessor_A<a>`". For the LCC card, this string identifies the frame in the format "`frame_F<f>`".
- `state`: A string that represents the current state of the node card. Values include:
  - `inServiceMode`
  - `unreachable`
  - `noEthernet`
  - `noCAN`
  - `online`
- `partNum`: A string that represents the part number of the node card.
- `sn`: A string that represents the serial number of the node card.
- `Bar code`: A 22-character string that represents the exact bar code label found on the node card.
- `ec`: An integer that represents the engineering change (EC) number of the node card.
- `firmware`: A string that represents the current firmware level of the node card.

- `cfBarcode`: A string that represents the value read from the serial number bar code of the compact flash card that is installed on the LCC. If this card is not an LCC card, this value is null. The format of the bar code matches that read off the compact flash card:
  - `SMART`
  - `Virtium`
- `cfPartNum`: A string that represents the part number of the compact flash card that is installed on the LCC. If this card is not an LCC, or the CF card does not support reporting this value, the value is null.
- `cfVendor`: A string that represents the name of the vendor of the compact flash card that is installed on the LCC. If this card is not an LCC card, this value is null. The other two supported values include:
  - `SMART`
  - `Virtium`
- `reportingLCC`: A string that represents whether this LCC is reporting the information for this **`GET`** command. This string returns **“**`yes`**”** if this LCC is reporting the status, **“**`no`**”** if this LCC’s status is being reported by another LCC, or null if this is not an LCC card.
- `primaryLCC`: A string that represents whether this LCC is acting as the primary LCC or the secondary LCC. This string returns **“**`yes`**”** if this LCC is the primary, **“**`no`**”** if this LCC is the secondary, or null if this is not an LCC card.

Reset node card

Run a reset on a node card.

Use the **`POST /v1/nodeCards/<ID>/reset`** command to run a reset on a node card.

Path parameter

The `ID` parameter is a decimal integer that is the unique identifier of the node card in the library.

When the **`POST /v1/nodeCards/65/reset`** command runs successfully, the system returns output to your display, as shown in Example 6-67.

Example 6-67 POST /v1/nodeCards/65/reset command

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros POST /v1/nodeCards/65/reset
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 0
Exit with code: 0
```

Node 65 is the LCC node card in frame 1. Resetting this node card causes the LCC to restart and post an event, as shown in the following examples:

- `Error Code: 2384`
- `Description: An LCC card was not responding over Ethernet`

The TS4500 features node cards installed as shown in the following order:

- NodeIdx 17 (x'11') XYC Acc A, NodeIdx 18 (x'12') XYC Acc B
- NodeIdx 49 (x'31') ACC Acc A, NodeIdx 50 (x'32') ACC Acc B
- NodeIdx 65 (x'41') LCC Frame 1
- NodeIdx 66 (x'42') LCC Frame 2
- NodeIdx 66 (x'42') LCC Frame 3, and so on

Power supplies

Display the status of both power supplies in each L or D frame

Use the **GET /v1/powerSupplies** command to get power supply status information.

When the **GET /v1/powerSupplies** command runs successfully, the system returns output to your display, as shown in Example 6-68.

Example 6-68 GET /v1/powerSupplies command

```
[root@tron ITDT]# ./itdt -f /dev/sg8 ros GET /v1/powerSupplies
[
  {
    "location": "powerSupply_F1PSa",
    "state": "online"
  },
  {
    "location": "powerSupply_F1PSb",
    "state": "online"
  },
  {
    "location": "powerSupply_F2PSa",
    "state": "online"
  },
  {
    "location": "powerSupply_F2PSb",
    "state": "online"
  }
]
```

Attributes

The following attributes are available:

- `location`: The string location of the power that is supplying the format "`powerSupply_F<frame>PS<power supply A or B>`"; for example, "`powerSupply_F1PSa`".
- `state`: The current operation mode of the power supply as last read by the LCC node card (can be obsolete if LCC node card is unreachable). This attribute includes the following values:
  - `failed`

    The power supply is showing a failure or is unreachable.Verify that the power supply is plugged in and cabled correctly from the rear of the frame. If this check does not resolve the issue, contact IBM for support.
  - `online`

    The power supply is plugged in and supplying power to the library.
  - `inServiceMode` *(future)*

    The power supply is powered off in preparation for service.

Report

The reports contain usage history and other data for resources in the library. A report is available for the following components:

- `drives`
- `library`
- `accessors`
- `cartridgeActivity`

Get drives report

The drive report contains usage history of the drive, including mounts, cleans, host I/O, compression rate, error rates, and environmental data. These reports are taken at 1-hour intervals and kept for up to 1 year. However, you can offload only 1 week of data at a time per **`GET`** command to ensure that the query response size does not get too large. If you do not include query parameters to limit the time frame, all reports for the last week are returned:

```
GET/v1/reports/drives
GET/v1/reports/drives?{<“after=time” | “before=time” | “location”>}
```

Optional query parameters

The following optional query parameters are available:

- `after`: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`” or “`YYYY-MM-DDThh:mm:ss`”; for example, “`2020-03-17T13:02:00`”. If the time zone is omitted, the system’s current time zone is used.
- `before`: A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`” or “`YYYY-MM-DDThh:mm:ss`”; for example, “`2020-03-17T13:02:00`”. If the time zone is omitted, the system’s current time zone is used.

When the **`GET /v1/reports/drives`** command runs successfully, the system returns output to your display, as shown in Example 6-69.

Example 6-69 GET /v1/reports/drives command output

```
[root@TRON ITDT]# ./itdt -f /dev/sg4 ros GET /v1/reports/drives
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 1132414
[
  {
    "location": "drive_F1C2R1",
    "sn": "0000078PG29A",
    "time": "2020-03-18T06:05:00-0600",
    "duration": 3600,
    "mounts": 0,
    "cleans": 0,
    "dataReadByHosts": 0,
    "dataWrittenByHosts": 0,
    "dataWrittenToCartridges": 0,
    "errorsCorrectedRead": 0,
    "errorsCorrectedWrite": 0,
    "errorsUncorrected": 0,
```

```
    "temperatureAverage": 22.0,
    "temperatureMin": 22.0,
    "temperatureMax": 22.0,
    "humidityAverage": 54.0,
    "humidityMin": 22.0,
    "humidityMax": 54.0
  }
...
```

The following command attributes are available:

- `location`: A string that represents the unique location of the tape drive. This string also is the unique identifier for the drive. The format is “`drive_F<f>C<c>R<r>`” where `'f'` is the frame, `'c'` is the column, and `'r'` is the row in which the drive is installed.
- `sn`: A string that represents the serial number of the drive. This string also can be used as a unique identifier for the tape drive.
- `time`: The time at which this data entry was recorded. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`”; for example, “`2020-03-17T13:02:00-0700`”.
- `duration`: The number of seconds this data entry represents in seconds. This string is approximately 3600 seconds, or 1 hour, but can differ slightly because of run time. The `'time'` attribute occurs at the end of this duration.
- `mounts`: An integer that represents the number of mounts that are performed per hour on this drive over the duration of this measurement.
- `cleans`: An integer that represents the number of cleans that are performed per hour on this drive over the duration of this measurement.
- `dataReadByHost`: An integer that represents the number of MB per hour of uncompressed data read by the host from the drive over the duration of this measurement.
- `dataWrittenToCartridges`: An integer that represents the number of MB per hour of compressed data that is written to the cartridges that are mounted in the drive over the duration of this measurement. The average compression ratio can be found by dividing `dataWritenByHosts` by `dataWrittenToCartridge`.
- `errorsCorrectedRead`: An integer that represents the number of corrected read errors per hour on the drive over the duration of this measurement.
- `errorsCorrectedWrite`: An integer that represents the number of corrected write errors per hour on the drive over the duration of this measurement.
- `errorsUncorrected`: An integer that represents the number of uncorrected errors per hour on the drive over the duration of this measurement.
- `temperatureAverage`: A floating-point number that is accurate to one decimal place that represents the average temperature in Celsius of the drive over the duration of this measurement.
- `temperatureMin`: A floating-point number that is accurate to one decimal place that represents the minimum temperature in Celsius that is reached in the drive over the duration of this measurement.
- `temperatureMax`: A floating-point number that is accurate to one decimal place that represents the minimum temperature in Celsius that is reached in the drive over the duration of this measurement.
- `humidityAverage`: A floating-point number that is accurate to one decimal place that represents the average percent humidity in the drive over the duration of this measurement. This value is null if the drive does not support reporting humidity measurements.

- `humidityMin`: A floating-point number that is accurate to one decimal place that represents the minimum (driest) percent humidity that is reached in the drive over the duration of this measurement. This value is null if the drive does not support reporting humidity measurements.
- `humidityMax`: A floating-point number that is accurate to one decimal place that represents the highest (wettest) percent humidity that is reached in the drive over the duration of this measurement. This value is null if the drive does not support reporting humidity measurements.

Get library report

The library report contains usage history of the library including mounts, inserts, ejects, moves, host I/O, and environmental data. These reports are taken at 1-hour intervals and kept for 1 year. However, you can offload only 1 week of data at a time per **GET** command to ensure that the query response size does not get too large. If you do not include query parameters to limit the time frame, all reports for the last week are returned:

```
GET/v1/reports/library
GET/v1/reports/library?{<"after=time" | "before=time" | "location">}
```

Optional query parameters

The following optional parameters are available:

- `after`: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`"; for example, "`2020-03-17T13:02:00`". If the time zone is omitted, the system's current time zone is used.
- `before`: A string that represents the time before which events should be viewed. Any event with a time after this value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`"; for example, "2020-03-17T13:02:00". If the time zone is omitted, the system's current time zone is used.

When the **`GET /v1/reports/library`** command is run, the output is similar to the output that is shown in Example 6-70.

Example 6-70 Get /v1/reports/library command output

```
[root@TRON ITDT]# ./itdt -f /dev/sg75 ros GET /v1/reports/library
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 54264
[
  {
    "time": "2020-03-18T06:05:06-0600",
    "duration": 3600,
    "mounts": 56665,
    "imports": 458,
    "exports": 789,
    "moves": 332,
    "dataReadByHosts": 4568970,
    "dataWrittenByHosts": 4534690,
    "dataWrittenToCartridges": 7899870,
    "temperatureAverage" 24.0,
    "temperatureMin": 23.0,
    "temperatureMax": 25.0,
```

```
    "humidityAverage": 45.0,
    "humidityMin": 40.0,
    "humidityMax": 52.0
  }
...
```


The following command attributes are available:

- `time`: The time at which this data entry was recorded. Time format is in ISO 8601 format `of` "`YYYY-MM-DDThh:mm:ss+-hhmm`"; for example, "`2018-09-17T23:02:00-0700`".
- `duration`: The number of seconds this data entry represents in seconds. This value is approximately 3600 seconds, or 1 hour, but might differ slightly because of run time. The `'time'` attribute occurs at the end of this duration.
- `mounts`: An integer that represents the number of mounts that are performed per hour on all drives in this library over the duration of this measurement.
- `imports`: An integer that represents the number of cartridges that are added to a library per hour over the duration of this measurement. A cartridge import is not complete until the host sends a SCSI move media command to the library for this cartridge or it was manually assigned to a Logical Library
- `exports`: An integer that represents the number of cartridges that are moved to the I/O station from a slot per hour over the duration of this measurement. A cartridge export is complete after the cartridge is physically moved to the I/O station.
- `moves`: An integer that represents the number of cartridge movement actions, including the get and put by the gripper, per hour that are run by all accessors in the library over the duration of this measurement.
- `dataReadByHosts`: An integer that represents the number of MB per hour of uncompressed data that is read by the host from all drives in the library over the duration of this measurement.
- `dataWrittenByHosts`: An integer that represents the number of MB per hour of uncompressed data that is written by the host to all drives in the library over the duration of this measurement.
- `dataWrittenToCartridges`: An integer that represents the number of MB per hour of compressed data that is written to the cartridges by all drives in the library over the duration of this measurement. The average compression ratio can be found by dividing `dataWritenByHosts` by `dataWrittenToCartridge`.
- `temperatureAverage`: A floating-point that is accurate to one decimal place number that represents the average temperature in Celsius in all drives in the library over the duration of this measurement.
- `temperatureMin`: A floating-point that is accurate to one decimal place number that represents the minimum temperature that is reached in Celsius in all drives in the library over the duration of this measurement.
- `temperatureMax`: A floating-point number that is accurate to one decimal place that represents the minimum temperature that is reached in Celsius accurate in all drives in the library over the duration of this measurement.
- `humidityAverage`: A floating-point number that is accurate to one decimal place that represents the average percent humidity that is reached in all drives in the library over the duration of this measurement.
- `humidityMin`: A floating-point number that is accurate to one decimal place that represents the minimum (driest) percent humidity that is reached in all drives in the library over the duration of this measurement.

- `humidityMax`: A floating-point number that is accurate to one decimal place that represents the highest (wettest) percent humidity that is reached in all drives in the library over the duration of this measurement.

Get Accessor Report

The accessor report contains usage history of the accessor including pivots, scans, distance traveled, and puts/gets done by the two grippers. These reports are taken on intervals of 1 hour and kept for 1 year. However, you can offload only 1 week of data at a time per **GET** command to ensure that the query response size does not get too large. If you do not include query parameters to limit the time frame, all reports for the last week are returned:

```
GET/v1/reports/accessors
GET/v1/reports/accessors?{<"after=time" | "before=time" | "location">}
```

Optional query parameters

The following optional query parameters are available:

- `after`: A string that represents the time after which events should be viewed. Any event with a time before this value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`"; for example, "`2020-03-17T13:02:00`". If the time zone is omitted, the system's current time zone is used.
- `before`: A string that represents the time before which events should be viewed. Any event with a time after value is not returned. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`" or "`YYYY-MM-DDThh:mm:ss`". For example, "`2020-03-17T13:02:00`". If the time zone is omitted, the system's current time zone is used.

When the **GET/v1/reports/accessors** command runs successfully, the system returns output to your display as shown in Example 6-71.

Example 6-71 Get /v1/reports/accessors output

```
[root@TRON ITDT]# ./itdt -f /dev/sg15  ros GET /v1/reports/accessors
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 36792
[
  {
    "location": "accessor_Aa",
    "time": "2020-03-18T07:05:00-0600",
    "duration": 3600,
    "pivots": 33334,
    "barCodeScans": 12330,
    "travelX": 1231230,
    "travelY": 1277823,
    "getsGripper1": 3498,
    "putsGripper1": 3546,
    "getsGripper2": 3851,
    "putsGripper2": 3215
  }
...
```

The following command attributes are available:

- `location`: A string describing the unique identifier for the accessor with a value of "accessor_Aa" or "accessor_Ab".
- `time`: The time at which this data entry was recorded. Time format is in ISO 8601 format of "`YYYY-MM-DDThh:mm:ss+-hhmm`"; for example, "2020-03-18T13:02:00-0700".
- `duration`: The number of seconds this data entry represents in seconds. This value is approximately 3600 seconds, or 1 hour, but might differ slightly because of run time. The `'time'` attribute occurs at the end of this duration.
- `pivots`: An integer that represents the number of pivots per hour that are performed by this accessor over the duration of this measurement.
- `barCodeScans`: An integer that represents the number of bar code scans per hour that are performed by this accessor over the duration of this measurement.
- `travelX`: An integer that represents the number meters per hour of movement in the X (horizontal) direction that are performed by this accessor over the duration of this measurement.
- `travelY`: An integer that represents the number meters per hour of movement in the Y (vertical) direction that are performed by this accessor over the duration of this measurement.
- `getsGripper1`: An integer that represents the number cartridge get actions per hour that are performed by gripper 1 of this accessor over the duration of this measurement.
- `putsGripper1`: An integer that represents the number cartridge put actions per hour that are performed by gripper 1 of this accessor over the duration of this measurement.
- `getsGripper2`: An integer that represents the number cartridge get actions per hour that are performed by gripper 2 of this accessor over the duration of this measurement.
- `putsGripper2`: An integer that represents the number cartridge put actions per hour that are performed by gripper 2 of this accessor over the duration of this measurement.

Slots

The `GET` command returns all storage slots in the system. Default sorting is by location in numerical order that is based on Frame, Column, and then, Row. On a large system, this list can be large.

Use the **`GET /v1/slots`** command for more information about slot status.

When the **`GET /v1/slots`** command runs successfully, the system returns output to your display, as shown in Example 6-72.

Example 6-72 GET /v1/slots command (truncated)

```
[root@tron ITDT]#./itdt -f /dev/sg8 ros GET /v1/slots
{
    "location": "slot_F1C9R1",
    "state": "normal",
    "contents": [
      "JJL051JJ",
      "CD0347JA",
      "GH0141JB",
      "CE 211JJ"
    ],
    "puts": 93,
    "putRetries": 0,
```

```
    "getRetries": 16,
    "tiers": 4
    "state": "online"
  }
]
```

Attributes

The following attributes are available:

- Location

  The string location of the slot in the library, not including the tier; for example, "`slot_F3C4R8`".

  An example for a full populated deep slot is shown in Example 6-72 on page 474.

- State

  The current state of the slot and includes the following values:

  – `inServiceMode`

    It is not selected as a cartridge destination; however, cartridges can be moved from the slot.

  – `unusable` (*future*)

    Not in use because it is not usable for media move operations because of the elastic capacity that is being disabled.

  – `inaccessible` (*future*)

    The slot is not accessible by an accessor. This issue often occurs because an accessor in the `noMovementAllowed` state prevents access to this slot, or all accessors are in a state that is unable to access.

  – `normal`

    Slot is ready for use.

- Contents

  An array of strings that shows the Volsers of the cartridges that are contained in the slot. They are listed in order of lowest numbered tier to highest numbered tier. Any empty tier is listed as null.

  This attribute is reported as a single-item array for a tier 0 (single-deep, door-side) slot.

- Puts

  An integer that shows the number of times a gripper has inserts a cartridge into this slot, including all put actions, such as shuffle and unshuffle moves.

- putRetries

  An integer that shows the number of times a retry was required while a gripper was inserting a cartridge into this slot.

- getRetries

  An integer that shows the number of times a retry was required while a gripper was retrieving a cartridge from this slot.

- Tiers

  An integer that represents the number of tiers that exist within this slot.

Tasks

The **GET /v1/tasks** command is used to get task information or to perform a task.

Various library tasks can be performed by using the **GET /v1/tasks** or the **POST /v1/tasks** commands. The task commands feature the following format:

```
GET /v1/tasks
GET /v1/tasks/<ID>
POST /v1/tasks {"type": "task type", "location": "hardware component"}
```

Path parameter

The `ID` parameter is a decimal integer that is the unique identifier of the task in the library.

Request body parameters

The following `type` request body parameters are available:

- `inventoryTier0and1`
- `inventoryAllTiers`
- `startDriveService`
- `completeDriveService`
- `calibrateAccessor`
- `calibrateLibrary`
- `calibrateFrame`
- `startAccessorService`
- `completeAccessorService`
- `testDrive`

When the **GET /v1/tasks** command runs successfully, the system returns output to your display as shown in Example 6-73.

Example 6-73 GET /v1/tasks command output

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/tasks
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 21448
[
  {
    "ID": 317,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "completed",
    "startTime": "2019-11-12T13:06:22-0600",
    "lastUpdateTime": "2019-11-12T13:20:38-0600",
    "percentComplete": 100,
    "duration": 856,
    "user": "admin",
    "description": "Library Code Update",
    "volser": null
  },
{
    "ID": 142,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "completed",
    "startTime": "2019-08-22T14:59:56-0600",
    "lastUpdateTime": "2019-08-22T15:13:08-0600",
```

```
    "percentComplete": 100,
    "duration": 792,
    "user": "Goku",
    "description": "Library Code Update",
    "volser": null
  },
  {
    "ID": 140,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "completed",
    "startTime": "2019-08-15T09:21:32-0600",
    "lastUpdateTime": "2019-08-15T09:34:30-0600",
    "percentComplete": 100,
    "duration": 778,
    "user": "admin",
    "description": "Library Code Update",
    "volser": null
  }
]
```


```
Exit with code: 0
```

The following command attributes are available:

- ID: A decimal integer that is the unique identifier of the task in the library.
- type: A string that represents the task type to run. Values can include:
  - inventoryTier0and1
  - inventoryAllTiers
  - calibrateLibrary
  - calibrateFrame
  - calibrateAccessor
  - testDrive
  - startDriveService
  - completeDriveService
  - startAccesorService
  - completeAccesorService
- location: A string that represents the hardware component what the task is affecting. This string can be null if no hardware was associated with this task. Hardware components include:
  - accessor_A<“a”|”b”> (for example, “accessor_Ab”)
  - gripper_A<“a”|”b”>G<“1”|”2”> (for example, “gripper_AbG1”)
  - frame_F<f> (for example, “frame_F1”)
  - drive_F<f>C<c>R<r> (for example, “drive_F1C3R23”)
  - ethernetPort_F<f>P<p> (for example, “ethernetPort_F2Pa”)
- state: A string showing the completion status of the task. Values include:
  - inProgress: The task is in progress. If supported, the “percentComplete” Task attribute shows how far along the task is and the “duration” attribute shows how long the task was running.
  - completed: The task completed successfully.
  - failed: The task failed. To find the reason for the failure, check the event list.

  - `aborted`: An LCC failover caused the task to be canceled automatically.
  - `canceled`: A user manually canceled the task or the library powered off.
- `startTime`: The time the task was started. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`”; for example, “`2018-09-17T23:02:00-0700`”.
- `lastUpdatedTime`: The time this task’s state was last updated. Time format is in ISO 8601 format of “`YYYY-MM-DDThh:mm:ss+-hhmm`”; for example, “`2018-09-17T23:02:00-0700`”.
- `percentComplete`: A percentage value that is accurate to one decimal place that shows how close the task is to being complete. This value is an *estimate* of the time the task takes, *not* a guarantee. This value is null for tasks that do not report this.
- `duration`: An integer that represents the number of seconds this task is running so far.
- `user`: A string that shows the name of the user whose action started the task.
- `volser`: For event types of “`verifyMedia`”, this field shows the VOLSER being verified. For other event types, this value is null.

If you want to monitor the status of only one task, use the **GET /v1/tasks/<ID>** command.

When the **GET /v1/tasks/<ID>** command runs successfully, the system returns output to your display as shown in Example 6-74.

Example 6-74 GET /v1/tasks/140 command output

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/tasks/140
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 290
[
  {
    "ID": 140,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "completed",
    "startTime": "2019-08-15T09:21:32-0600",
    "lastUpdateTime": "2019-08-15T09:34:30-0600",
    "percentComplete": 100,
    "duration": 778,
    "user": "admin",
    "description": "Library Code Update",
    "volser": null
  }
]

Exit with code: 0
```

If you want to start a new task on the library, use the **POST /v1/tasks {“type”: “task type”, “location”: “hardware component”}** command.

When the **POST /v1/tasks {“type”: “task type”, “location”: “hardware component”}** command runs successfully, the system returns output to your display, as shown in Example 6-75 on page 479.

Example 6-75 POST /v1/tasks/ {"type": "inventoryTier0and1", "location": "library"} *command*

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros POST /v1/tasks/ {"type":
"inventoryTier0and1", "location": "library"}
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
{
  "ID": "131"
}

Exit with code: 0
```

After the task is created it can be monitored with the **`GET /v1/tasks`** command. Use the ID that is presented in the command output (in this case, ID: 131). In Example 6-76, we can see that task 131 is in progress and is 29% complete.

Example 6-76 GET /v1/tasks/131 command

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros GET /v1/tasks/131
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 299
[
  {
    "ID": "131",
    "type": "inventoryTier0and1",
    "location": "library",
    "state": "inProgress",
    "startTime": "2019-11-20T06:43:57-0600",
    "lastUpdateTime": "2019-11-20T06:45:21-0600",
    "percentComplete": "29",
    "duration": "84",
    "user": "SCSIControlPath_F1C3R4",
    "description": "Inventory",
    "volser": null
  }
]

Exit with code: 0
```

Starting service on a drive

Initiate a service action on a drive, which puts the drive in the "`in service`" state. This process starts a long running task in the library that is visible from the GUI:

`POST /v1/tasks {"type": "startDriveService", "location": "drive_F<f>C<c>R<r>"}`

Path parameters

The path parameters are available:

- `type`: A string that represents the task type to run; in this case, "`startDriveService`".
- `location`: A string that represents the location of the tape drive. This string also is the unique identifier for the drive. The format is "`drive_F<f>C<c>R<r>`" where "`f`" is the frame, "`c`" is the column, and "`r`" is the row in which the drive is installed.

When the **POST /v1/tasks {“type”: ”startDriveService”, “location”: “drive_F<f>C<c>R<r>”}** command runs successfully, the system returns output to your display, as shown in Example 6-77.

Example 6-77 POST /v1/tasks/ {"type": "startDriveService", "location": "drive_F1C2R2"}*command*

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros POST /v1/tasks/ {"type":
"startDriveService", "location": "drive_F1C2R2"}
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
{
  "ID": "129"
}

Exit with code: 0
```

You can also see that task ID: 129 was created. This cancellation can be monitored by using the **GET /v1/tasks** command.

Complete service on a drive

Complete a service action on a drive. This process takes the drive out of the “in service” state. It also starts a long running task in the library that is visible from the GUI:

```
POST /v1/tasks [{“type”: “completeDriveService”, “location”:
“drive_F<f>C<c>R<r>”}]
```

Path parameters

The following path parameters are available:

- type: A string that represents the task type to run; in this case “completeDriveService”.
- location: A string that represents the location of the tape drive. This string also is the unique identifier for the drive. The format is “drive_F<f>C<c>R<r>” where “f” is the frame, “c” is the column, and “r” is the row in which the drive is installed.

When the **POST /v1/tasks {“type”: “completeDriveService”, “location”: “drive_F<f>C<c>R<r>”}** command runs successfully, the system returns output to your display, as shown in Example 6-78.

Example 6-78 POST /v1/tasks/ {"type": "completeDriveService", "location": "drive_F1C2R2"} *command*

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros POST /v1/tasks/ {"type":
"completeDriveService", "location": "drive_F1C2R2"}
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
{
  "ID": "133"
}

Exit with code: 0
```

You can also see that task ID: 133 was created. This can be monitored by using the **GET /v1/tasks** command.

The drive status changes from “`offline`” to “`disconnect`”, then back to “`offline`” while running a drive test. After the drive test is complete, the diagnostic cartridge is unloaded and the drive status changes back to “`online`”.

Calibrate accessor

Run a calibration to an accessor. This process starts a long-running task in the library that is visible from the GUI:

`POST /v1/tasks {“type”: ”calibrateAccessor”, “location”: ”accessor_A<”a”|”b”>”}`

Path parameters

The following path parameters are available:

- `type`: A string that represents the task type to run; in this case, “`calibrateAccessor`”.
- `location`: A string that describes the unique identifier for the accessor the features a format of: “`accessor_A<“a”|“b”>`”.

When the **`POST /v1/tasks {“type”: “calibrateAccessor”, “location”: “accessor_A<“a”|“b”>”}`** command runs successfully, the system returns output to your display, as shown in Example 6-79.

Example 6-79 POST /v1/tasks/ {“type”: “calibrateAccessor, “location”: "accessor_Aa"} command

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros POST /v1/tasks {"type":
"calibrateAccessorService", "location": "accessor_Aa"}
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
{
  "ID": "135"
}

Exit with code: 0
```

You can also see that task ID: 135 was created. This process can be monitored by using the **`GET /v1/tasks`**, command.

Start service on accessor

Start a service action on an accessor, which places the accessor in the “`in service`” state. This process starts a long-running task in the library that is visible from the GUI:

`POST /v1/tasks {“type”: “startAccessorService”, “location”: “accessor_A<“a”|“b”>”}`

Path parameters

The following path parameters are available:

- `type`: A string that represents the task type to run; in this case “`startAccessorService`”.
- `location`: A string that describes the unique identifier for the accessor with a format of “`accessor_A<“a”|“b”>`”.

When the **`POST /v1/tasks {“type”: “startAccessorService”, “location”: “accessor_A<“a”|“b”>”}`** command runs successfully, the system returns output to your display, as shown in Example 6-80 on page 482.

Example 6-80 POST /v1/tasks {"type": "startAccessorService", "location": "accessor_Aa"} command

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros POST /v1/tasks {"type":
"startAccessorService", "location": "accessor_Aa"}
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
{
  "ID": "137"
}

Exit with code: 0
```

You can also see that task ID: 137 was created. This process can be monitored by using the **GET /v1/tasks** command.

Complete service on accessor

Complete a service action on an accessor. This process takes the accessor out of the "in service" state. It also starts a long-running task in the library that is visible from the GUI:

```
POST /v1/tasks {"type": "completeAccessorService", "location": "accessor_A<"a"|
"b">"}
```

Path parameters

The following path parameters are available:

- type: A string that represents the task type to run; in this case "completeAccessorService".
- location: A string that describes the unique identifier for the accessor with a format of "accessor_A<"a"|"b">".

When the **POST /v1/tasks {"type": "completeAccessorService", "location": "accessor_A<"a"|"b">"} c**ommand runs successfully, the system returns output to your display, as shown in Example 6-81.

Example 6-81 POST /v1/tasks {"type": "completeAccessorService", "location": "accessor_Aa"}

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros POST /v1/tasks {"type":
"completeAccessorService", "location": "accessor_Aa"}
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
{
  "ID": "139"
}
Exit with code: 0
```

You can also see that task ID: 139 was created. This process can be monitored by using the **GET /v1/tasks** command.

The following steps are performed by the accessor during the Complete Service operation:

1. FRU Complete ACC-A - Rezero Request
2. FRU Complete ACC-A - Calibrate Accessor
3. FRU Complete ACC-A - Calibrate ISB (Internal Service Bay)
4. FRU Complete ACC-A - Verify Library

5. FRU Complete ACC-A - Issue Final Rezero
6. FRU Complete ACC-A - Done

Test drive

Start a test of a drive operation by using a diagnostic cartridge. This process starts a long-running task in the library that is visible from the GUI:

```
POST /v1/tasks/ '{"type":"testDrive","location":"drive_F<f>C<c>R<r>"}'
```

Path parameters

The following path parameters are available:

- `type`: A string that represents the task type to run; in this case "`testDrive`".
- `location`: A string that represents the location of the tape drive. This string also is the unique identifier for the drive. The format is "`drive_F<f>C<c>R<r>`" where "`f`" is the frame, "`c`" is the column, and "`r`" is the row in which the drive is installed.

When the **`POST /v1/tasks/ '{"type":"testDrive","location":"drive_F<f>C<c>R<r>"}'`** command runs successfully, the system returns output to your display, as shown in Example 6-82.

Example 6-82 POST /v1/tasks/ '{"type": "testDrive, "location":"drive_F1C2R2"}' command

```
[root@TRON ITDT]# ./itdt -f /dev/sg43 ros POST /v1/tasks/
'{"type":"testDrive,"location":"drive_F2C2R2"}'
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
{
  "ID": "145"
}
Exit with code: 0
```

You can also see that task ID: 145 was created. This process can be monitored by using **`GET /v1/tasks`** command.

Cartridge Movement Request - Move To Drive

To request movement of a cartridge, a work item is created in the work queue. The ID of the work item is returned after it is submitted:

```
POST /v1/workItems [{"type": "moveToDrive", "cartridge":
<volser>,"sourceInternalAddress": <internalAddress>, "destinationLocation":
<location>, "destinationSN": < serialNumber >}]
```

Request body parameters

The following path parameters are available:

- `type`: A string that represents the type of work item that is created. In this case, `moveToDrive`.
- `cartridge`: A string that represents the VOLSER of the cartridge to be moved. If duplicate VOLSERs exist for this cartridge in the library, use `sourceInternalAddress` instead.
- `sourceInternalAddress`: A string that represents the internal address of the location of the cartridge that is to be moved. The `sourceInternalAddress` string might change if the cartridge is assigned or unassigned from a logical library or if the cartridge is moved by the host or system. This value must be queried immediately before use.

- `destinationLocation`: A string that represents the location of the destination tape drive (for example, "`drive_F2C4R2`").
- `destinationSN`: A string that represents the location of the destination tape drive (for example, "`00078B3455`").

When the **`POST /v1/workItems [{"type": "moveToDrive", "cartridge": <volser>,"sourceInternalAddress": <internalAddress>, "destinationLocation": <location>, "destinationSN": < serialNumber >}]`** command runs successfully, the system returns output to your display, as shown in Example 6-83.

Example 6-83 Returned output

```
[root@tron ITDT]# ./itdt -f /dev/sg8 ros POST /v1/workItems '{"type":
"moveToDrive", "cartridge": "F20223L9", "destinationLocation": "drive_F2C2R3"}'
{
  "ID": "0"
}
```

This process can be monitored by using the following commands:

- `./itdt -f /dev/sg8 ros GET /v1/drives/drive_F2C2R3`
- `./itdt -f /dev/sg8 ros GET /v1/dataCartridges/F20224L9`

The **`drives`** and **`dataCartridges`** commands feature output that is similar to the following examples:

```
"volser": "F20223L9",
"state": "normal",
"accessible": "normal",
"location": "drive_F2C2R3",
```

> **Note:** Cartridge and drive has to be in the same logical library.

Cartridge Movement Request - Move To IO Station

To request movement of a cartridge, a work item is created in the work queue. The ID of the work item is returned after it is submitted. It moves a cartridge to an I/O station that is chosen by the library:

```
POST /v1/workItems [{"type": "moveToIOStation", "cartridge": <volser>,
"sourceInternalAddress": <internalAddress>}]
```

Request body parameters

The following path parameters are available:

- `type`: A string that represents the type of work item that is being created (in this case, `moveToIOStation`).
- `cartridge`: A string that represents the VOLSER of the cartridge that is to be moved. If duplicate VOLSERs exist for this cartridge in the library, the `sourceInternalAddress` must be used instead.
- `sourceInternalAddress`: A string that represents the internal address of the location of the cartridge to be moved. This string might change if the cartridge is assigned or unassigned from a logical library or if the cartridge is moved by the host or system. This value must be queried immediately before use.

When the **POST /v1/workItems [{"type": "moveToIOStation", "cartridge": <volser>, "sourceInternalAddress": <internalAddress>}]** command runs successfully, the system returns output to your display, as shown in Example 6-84 on page 485.

Example 6-84 Returned output

```
./itdt -f /dev/sg8 ros POST /v1/workItems '{"type": "moveToIOStation",
"cartridge": "SGP396L3"}'
}

Note:
The source must be specified by either entering the cartridge or
sourceInternalAddress. Only one of these sources must be specified.

Note: In case there is no magazine in the I/O Station, or the magazine is already
full, the cartridge will not physically move but the internal address will be
changed to any virtual IO (VIO) slot number. (i.e. 0x301).

This process can be monitored by using the following command:
./itdt -f /dev/sg8 ros GET /v1/dataCartridges/SGP396L3

Will give you analogous the following information:
"volser": "SGP396L3",
"state": "normal",
"accessible": "normal",
"location": "ioSlot_F1IOuR1",
"mediaType": "LTO",
"encrypted": null,
"mostRecentVerification": null,
"mostRecentUsage": "2021-10-12T05:07:39-0700",
"logicalLibrary": null,
"elementAddress": null,
"internalAddress": "FF0302"
```

Cartridge Movement Request - Move Cartridge To Slot

To request movement of a cartridge, a work item is created in the work queue. The ID of the work item is returned after it is submitted. It moves a cartridge to a slot that is chosen by the library:

```
POST /v1/workItems [{"type": "moveToSlot", "cartridge":
<volser>,"sourceInternalAddress": <internalAddress>}]
```

Request body parameters

The following path parameters are available:

- `type`: A string that represents the type of work item that is being created (in this case, `moveToSlot`).
- `cartridge`: A string that represents the VOLSER of the cartridge to be moved. If duplicate VOLSERs exist for this cartridge in the library, the `sourceInternalAddress` must be used instead.
- `sourceInternalAddress:` A string that represents the internal address of the location of the cartridge that is to be moved. This string might change if the cartridge is assigned or unassigned from a logical library or if the cartridge is moved by the host or system. This value must be queried immediately before use.

- `type`: A string that represents the type of work item that is created (in this case, `moveToDrive`).

When the **POST /v1/workItems [{"type": "moveToSlot", "cartridge": <volser>,"sourceInternalAddress": <internalAddress>}]** command runs successfully, the system returns output to your display, as shown in Example 6-85 on page 486.

Example 6-85 Returned output

```
./itdt -f /dev/sg8 ros POST /v1/workItems '{"type": "moveToSlot", "cartridge":
"SGP396L3"}'
{
  "ID": "0"
}
```

> **Note:** Consider the following points:
>
> - The source must be specified by entering the cartridge or `sourceInternalAddress`. Only one of these sources must be specified.
> - As of this writing, only physical transport is available between the I/O station and storage slot. If an attempt is made to move a cassette between the storage slots, only the internal address is changed, not the location. A target address is specified in a later release.

This process can be monitored by using the following command:

```
./itdt -f /dev/sg8 ros GET /v1/dataCartridges/SGP396L3
```

The **dataCartridges** command features output that is similar to the following example:

```
"volser": "SGP396L3",
"state": "normal",
"accessible": "normal",
"location": "slot_F2C2R18T0",
"mediaType": "LTO",
"encrypted": null,
"mostRecentVerification": null,
"mostRecentUsage": "2021-10-12T05:09:40-0700",
"logicalLibrary": null,
"elementAddress": null,
"internalAddress": "FF040B"
```

6.4 HTTP return codes

HTTP defines 40 standard status codes that can be used to convey the results of a user's request. The status codes are divided into the following categories:

- 1xx: Informational: Communicates transfer protocol-level information.
- 2xx: Success: Indicates that the user's request was accepted successfully.
- 3xx: Redirection: Indicates that the user must take some other action to complete their request.
- 4xx: Client Error: This category of error status codes indicate a possible user error.
- 5xx: Server Error: The server takes responsibility for these error status codes.

Examples

As shown in Example 6-86, return code 201 confirms that the requested task was accepted.

Example 6-86 HTTP/1.1 201 Return code

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros POST /v1/tasks/ {"type":
"inventoryTier0and1", "location": "library"}
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 13
Exit with code: 0
```

As shown in Example 6-87, return code 404 states that an error exists in the command.

Example 6-87 HTTP/1.1 404 Return code

```
[root@TRON ITDT]# ./itdt -f /dev/sg34 ros GET /v1/tasks/325
HTTP/1.1 404 Not Found
Content-Type: application/json
Content-Length: 88
Exit with code: 0
```

Related publications

The publications that are listed in this section are particularly suitable for a more detailed discussion of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- *IBM Tape Library Guide for Open Systems*, SG24-5946
- *IBM Spectrum Archive Enterprise Edition V1.3.2.2: Installation and Configuration Guide, SG24-8333*

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials, at the following website:

ibm.com/redbooks

Other publications

The following publications are also relevant as further information sources:

- *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130:

  https://www.ibm.com/support/pages/ibm-tape-device-drivers-installation-and-users-guide
- *IBM TotalStorage Enterprise Tape System 3592 Introduction and Planning Guide,* GA32-0555:

  https://www.ibm.com/support/pages/ibm-system-storage-ts1120-tape-drive-and-controller-introduction-and-planning-guide-3592-models-j1a-e05-e06-eu6-e07-e08-and-j70-c06-c07-controllers
- *IBM TotalStorage LTO Ultrium Tape Drive SCSI Reference (LTO-5 through LTO-9)*, GA32-0928:

  https://www.ibm.com/support/pages/ibm%C2%AE-totalstorage%C2%AE-lto-ultrium-tape-drive-scsi-reference-lto-5-through-lto-9
- *IBM 3592 SCSI Reference*, GA32-0562:

  https://www.ibm.com/support/pages/ibm-system-storage-tape-drive-3592-scsi-reference
- *IBM Tape System Library Manager User's Guide*, GA32-2208:

  https://www.ibm.com/resources/publications/OutputPubsDetails?PubID=GA32-2208-06

© Copyright IBM Corp. 2014, 2021, 2022. All rights reserved.

Online resources

The following web pages are helpful for more information:

- IBM System Storage Interoperation Center (SSIC):

  https://www.ibm.com/systems/support/storage/ssic/interoperability.wss
- IBM Spectrum Archive Library Edition (LE) documentation in IBM Documentation:

  https://www.ibm.com/docs/en/spectrum-archive-le
- IBM Spectrum Archive Enterprise Edition documentation in IBM Documentation:

  https://www.ibm.com/docs/en/spectrum-archive-ee
- Library and drive firmware packages in IBM Fix Central:

  http://www.ibm.com/support/fixcentral
- IBM Security Guardium Key Lifecycle Manager in IBM Documentation:

  https://www.ibm.com/docs/en/sgklm
- IBM Security Guardium Key Lifecycle Manager product ordering and resources website:

  https://www.ibm.com/products/ibm-security-key-lifecycle-manager
- IBM tape storage products for data protection and long-term retention:

  https://www.ibm.com/storage/tape
- IBM Linear Tape File System (LTFS):

  https://www.ibm.com/docs/en/spectrum-archive-ee/1.2.6.0?topic=overview-ltfs-format
- TS4500 documentation:

  https://www.ibm.com/docs/en/ts4500-tape-library
- ISV matrix, which is available to download as a PDF file from the "Independent Software Vendor Matrix (ISV) for IBM TotalStorage 3592 tape drives and LTO" from the TS4500 product web page under Resources link:

  https://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=WH&infotype=SA&appname=STGE_IV_IV_USEN&htmlfid=IVL12347USEN&attachment=IVL12347USEN.PDF
- IBM Data Protection and Retention (DP&R) System Connectivity and Security:

  https://www.ibm.com/support/pages/ibm-dri-system-connectivity-and-security-v332

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM TS4500 R8 Tape Library Guide

SG24-8235-09

ISBN 0738460419

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

SG24-8235-09

ISBN 0738460419

Printed in U.S.A.

Get connected

Redbooks®

ibm.com/redbooks