

IBM Diamondback Tape Library Guide

Phil Gerrard

Jesus Cabrera

Dylan Carlson

Jesus Cervantes

Fabian Corona

Stephen Solewin

Erwin Zwemmer



 **Cloud**

Storage



IBM Redbooks

IBM Diamondback Tape Library Guide

June 2024

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

First Edition (June 2024)

This edition applies to Version 2.9.2 of the IBM Diamondback tape library

© Copyright International Business Machines Corporation 2024. All rights reserved.

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

Contents

Notices	ix
Trademarks	x
Preface	xi
Now you can become a published author, too!	xiii
Comments welcome.	xiii
Stay connected to IBM Redbooks	xiii
Chapter 1. IBM Diamondback tape library	11
1.1 IBM Diamondback product description	12
1.2 IBM Diamondback library components	13
1.3 Feature codes for IBM Diamondback	17
1.3.1 External IBM Tape System Service Console (TSSC)	17
1.3.2 Top of rack frame	18
1.4 Host platforms and device drivers	19
1.5 IBM Diamondback specifications	19
1.5.1 Physical specifications	19
1.5.2 Floor space requirements	19
1.5.3 Power and cooling	21
Chapter 2. IBM Diamondback Linear Tape-Open	33
2.1 IBM LTO tape drives for the Diamondback library	34
2.1.1 LTO core technology	34
2.1.2 Data compression	36
2.1.3 LTO tape media	37
2.1.4 LTO 9 media optimization	42
2.1.5 Physical attachment	43
2.1.6 Performance specifications	46
2.2 IBM LTO 9 tape drives	46
2.2.1 LTO 9 tape drive features	47
2.2.2 Compatibility	47
2.2.3 Connectivity	48
2.2.4 Performance	48
2.2.5 Dynamic speed matching	49
2.2.6 Encryption	49
2.2.7 IBM Storage Archive software application	49
2.2.8 Surface Control Guiding Mechanism	50
2.2.9 Statistical Analysis and Reporting System (SARS)	50
2.2.10 Media partitioning	50
2.2.11 IBM Storage Archive software application	51
Chapter 3. IBM Diamondback Tape Library Supported Software	33
3.1 IBM Storage Archive	34
3.1.1 Key capabilities	34
3.1.2 Benefits	34
3.2 Linear Tape File System	35
3.2.1 LTFS partitioning support	36
3.2.2 Wrap-wise partitioning	37
3.2.3 IBM Storage Archive Editions	38

Chapter 4. Security and encryption	49
4.1 Tape encryption overview	50
4.1.1 Encryption-enabled tape drives	50
4.1.2 Encryption key management	50
4.2 Encryption policy	51
4.2.1 Encryption methods	51
4.2.2 Encryption key management on the Diamondback	52
4.2.3 Application-managed encryption	53
4.2.4 Library-managed encryption	54
4.2.5 Prerequisites for using encryption on the IBM Diamondback tape library	54
4.3 Additional Security options	57
4.3.1 Air gap	57
4.3.2 Safeguarded tape	57
4.3.3 WORM tape	57
Chapter 5. Diamondback management graphical user interface (GUI)	53
5.1 Using the Diamondback management GUI	54
5.1.1 Connecting to the management GUI	54
5.1.2 System summary window	56
5.2 Monitoring	58
5.2.1 Library	58
5.2.2 Actions and remote management of the library	62
5.2.3 Events	69
5.3 Library menu	74
5.3.1 Creating a logical library	81
5.3.2 Slots	90
5.4 Drives menu	91
5.4.1 Drives	91
5.4.2 Drives by Logical Library page	97
5.4.3 Fibre Channel Ports option	98
5.4.4 SAS Ports	101
5.5 Cartridges menu	103
5.5.1 Cartridges	103
5.5.2 Cartridges by Logical Library page	111
5.5.3 VOLSER ranges	115
5.5.4 VOLSER range by logical library	117
5.6 Access menu	118
5.6.1 Users	118
5.6.2 Roles	125
5.7 Settings	127
5.7.1 Library	127
5.7.2 Networking	129
5.7.3 Notifications	132
5.7.4 GUI Preferences page	150
5.8 Remote support	151
5.8.1 Remote support by using an external TSSC	154
5.8.2 Remote support security	154
5.8.3 Saving and restoring the Diamondback Backup to the TSSC	158
5.9 IBM Net Promoter Score and Feedback	162
Chapter 6. IBM Diamondback REST API	57
6.1 REST API Overview	58
6.1.1 Commands	58
6.1.2 Endpoints	58

6.1.3 Parameters	104
6.2 Using sg3_utils (Linux)	105
6.2.1 Example	105
6.2.2 Endpoints	106
6.3 Using ITDT for RoS	106
6.3.1 Examples	107
6.3.2 Endpoints and ITDT	108
6.4 REST over Ethernet	108
6.4.1 Authentication	108
6.4.2 Examples	109
6.4.3 Endpoints	110
6.5 HTTP return codes	110
Chapter 7. Diamondback self-support	121
7.1 Diamondback self-support and serviceability	122
7.1.1 Self-support and CRUs	122
7.1.2 QR codes	123
7.2 Self-support scenario and issue resolution example	125
Appendix A. Feature codes for IBM Diamondback	129
Related publications and links	139
Service Video Library	139
IBM Redbooks	139
Publications	139

Notices

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

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

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

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

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

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <https://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®	IBM Z®	Redbooks (logo)  ®
Guardium®	NPS®	Tivoli®
IBM®	ProtecTIER®	XIV®
IBM Security®	RACF®	z/OS®
IBM Spectrum®	Redbooks®	

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, 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® Diamondback tape library is the next-generation, ultra-high-density tape storage solution available from IBM. The Diamondback is designed to easily scale out and fit into any data center or server room, with each frame and has the footprint of only a standard server rack. Designed with a focus on simplicity and density, the library is designed to be deployed quickly and easily managed by using either a REST API or its browser-based GUI. Maintenance and support are also simplified by using Customer Replaceable Units (CRUs). CRUs and a built-in call home can greatly reduce the number of parts to be kept on hand and reduce downtime in the event of a failure.

This next-generation tape library is designed to help midsized and large enterprises meet modern cloud storage challenges by simplifying data migration and administration. The challenges include increasing data volumes, rising storage footprint costs, and reducing the complexity of IT training and system management.

This IBM Redbooks® publication provides a detailed overview of the new Diamondback library, which includes the following features:

- ▶ Standard server rack dimensions based on Open Compute Project racks, installs into data centers, and is serviced like a server rack
- ▶ Maximum storage capacity of 27 PB with LTO 9 cartridges and a maximum of 69 PB with 2.5:1 compression, a 2.7x improvement over LTO with TS4500 single frame
- ▶ Reduce installation time with the ability to be delivered fully loaded with tape cartridges
- ▶ Optimize storage with Ultra high density (UHD) storage slots
- ▶ Automate management through a REST API
- ▶ Simplified service by using CRUs reduces downtime and number of parts to keep on hand
- ▶ Automatic control-path and data-path fail-over
- ▶ Tape-drive encryption
- ▶ WORM media support
- ▶ Persistent worldwide name
- ▶ Multipath architecture
- ▶ Drive and media exception reporting
- ▶ Remote drive and media management
- ▶ Advanced Library Management System (ALMS)

Also, this book describes detailed information about the IBM Diamondback components, feature codes, specifications, supported tape drives, encryption, system management, security features, and REST features.

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

Authors

This book was produced by a while they worked at IBM Redbooks, Tucson Center.

Phillip Gerrard is a Project Leader for the International Technical Support Organization working out of Beaverton, Oregon. As part of IBM for over 15 years, he has authored and contributed to hundreds of technical documents published to IBM.com and worked directly with IBM's largest customers to resolve critical situations. As a team lead and Subject Matter Expert for the IBM Spectrum® Protect support team, he is experienced in leading and growing international teams of talented IBM employees, developing and implementing team processes, and creating and delivering education. Phillip holds a degree in computer science and business administration from Oregon State University.

José Jesús Cabrera Vega is a Senior Test Engineer in IBM México Software Lab. He has a bachelor's degree in computer science and joined IBM's System Level Test Open Tape team in 2015. As the team lead of the SLT Open Tape team since 2019, he is responsible for leading Diamondback, TS4500, TS3500 library testing as well as maintaining and updating the interoperability roadmap.

Dylan Carlson is an IBM Tape Library Firmware Developer and Field Analysis Lead with the IBM Systems Tape and Tape Automation development teams. As a REST over Ethernet Innovator and TS4500 REST interface designer, he has helped propel tape to new heights of reliability, efficiency, and sustainability. As a key member of the tape development team, he is focused on revolutionizing the way devices communicate and interact in the digital realm.

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 has been with IBM since 2012. As a certified ISTQB FL Tester, for the past 9 years he has been engaged with IBM's next generation of storage management systems, the IBM TS4500 Tape Library in the Product Test area (focused in Functional Verification Test) and has served as the team leader for the last 5 years.

Fabian Corona 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 is also a certified ISTQB Advanced Level Test Manager.

Stephen Solewin has 34 years of experience in the IT industry, most recently at Kyndryl, and 25 years with IBM. He is a world-renown expert on IBM Flash Systems, A9000, and XIV®. He has experience with many aspects of the IBM storage portfolio, including virtualization, disk, tape, and storage area networks. He is now a Global Storage Architect, working to enable storage technology for Kyndryl customers worldwide, creating best practices, and working with customers to ensure that their storage environments are implemented correctly and finely tuned. Steve attended the University of Arizona, obtaining a BS degree in Electrical Engineering, with Honors.

Erwin Zwemmer is an IBM Remote Technical Support (RTS) Specialist and works as a Subject Matter Expert (SME) for IBM EMEA Tape Storage, and takes advantage of the close collaboration with product development, support centers, and other business units to provide seamless problem resolution and drive continuous quality improvements. In 1995, Erwin joined IBM Netherlands where he worked as a Customer Engineer for AS/400 systems and multivendor products. In 2001, he moved to IBM Germany where he continues to support customers in EMEA with complex TS4500 installation and on-demand support on site. He has been a contributing author of several IBM Redbooks.

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

Shawn Nave, Pete Wendler, Lee Jesionowski
IBM Tucson Center

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 IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, IBM Redbooks
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

<https://www.linkedin.com/groups/2130806>

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

<https://www.redbooks.ibm.com/subscribe>

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

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



IBM Diamondback tape library

Building on more than two decades of storage innovation, the IBM Diamondback tape library (Diamondback) provides ultra high-density long-term data storage system within a standard rack footprint. The Diamondback is the newest addition to IBM's family of tape libraries. IBM Diamondback provides the important air gap cyber resiliency that only tape can provide. Write once read many (WORM) media, data encryption, and Safeguarded Tape features provide additional layers of protection for all of your important data. Designed with serviceability and sustainability in mind, Diamondback provides a very dense enterprise tape library. IBM Diamondback can be deployed as a scaled-out solution with 1548 usable slots for LTO 9 cartridges, and a maximum of 14 LTO 9 drives. Diamondback provides 27.8 PB of uncompressed capacity per library or 69.7 PB with a 2.5:1 compression ratio. The Diamondback library provides the highest data density available from IBM, and is over 2 times more dense than the IBM TS4500. Diamondback comes with a GUI and a full REST API, which can make managing one or more Diamondback libraries easier. LTO 9 drives are available with an 8 Gbps multi-mode Fibre Channel interface or 12 Gbps SAS interface.

IBM Diamondback comes in a standard 42U 19-inch rack, which can make it easier to plan, place and incorporate the tape library into a new or an existing data storage location. Diamondback provides an optional top of frame rack with 5U of additional space, taking the frame to a total height of 48U.

This chapter includes the following topics:

- ▶ 1.1, "IBM Diamondback product description" on page 12
- ▶ 1.2, "IBM Diamondback library components" on page 13
- ▶ 1.3, "Feature codes for IBM Diamondback" on page 17
- ▶ 1.4, "Host platforms and device drivers" on page 19
- ▶ 1.5, "IBM Diamondback specifications" on page 19

1.1 IBM Diamondback product description

The IBM Diamondback library is a single frame, stand-alone tape library that provides very high density, high performance, automated tape handling. Figure 1-1 shows a Diamondback with the optional top of frame rack added.



Figure 1-1 IBM Diamondback tape library with optional top of frame rack.

The IBM Diamondback library is designed to be serviced by the customer, with a few Customer Replaceable Units (CRUs) and online videos showing the simple repair procedures. For more information, see 7.1, “Diamondback self-support and serviceability” on page 122. Most parts, including the cartridge accessor assembly, can be replaced in 2 minutes or less. Diamondback is also available with full IBM Service provided by IBM, with both Basic and Premium levels of service available. Diamondback can be ordered and shipped with media preloaded. Each cartridge slot is an ultrahigh density slot, with the capability to hold up to 9 LTO 9 tape cartridges. If a library is ordered with a number of cartridges not evenly divisible by nine, the additional cartridges (somewhere between 1 and 8) are included and packaged inside the bottom of the library.

The Diamondback library supports a single robotic accessor, which includes dual redundant grippers, axis motors and a bar code scanner in one CRU.

The IBM Diamondback library includes the following capabilities and features:

- ▶ Most of the parts are Customer replaceable units (CRU)
- ▶ 1584 total slots, with 1548 usable
- ▶ Support for an external IBM Total Storage System Console (TSSC)
- ▶ Web-based user interface
- ▶ 10 cartridge service magazine
- ▶ Up to four library-managed encryption (LME) key paths (or key servers) per logical library
- ▶ A maximum of fourteen LTO 9 drives with two available drive types that cannot be mixed:
 - feature code AGH1 for 8 Gbps dual-port multimode Fibre Channel
 - feature code AGH3 for 12 Gbps dual port SAS
- ▶ LDAP support
- ▶ Safeguarded Tape
- ▶ Remote monitoring by using syslog
- ▶ Dual grippers
- ▶ Optional preloaded media inventory

For a full list of features of the Diamondback Tape Library, see [Features](#).

1.2 IBM Diamondback library components

The IBM Diamondback library consists of a tape library frame, with an optional top rack frame. Figure 1-2 on page 14 shows the main components visible from the front of the library and figure Figure 1-3 on page 15 shows the rear view.




Figure 1-2 Front internal view of the IBM Diamondback library with media restraint brackets in place

Table 1-1 BM Diamondback internal components front view in Figure 1-2

Number	Component
1	Library frame
2	Vertical robotic accessor rails
3	Cartridge accessor assembly
4	IBM LTO 9 tape drives
5	Cartridge storage slots
6	Front door
7	Front door safety switch
8	Door sensor




Figure 1-3 Rear internal view of the IBM Diamondback library

Table 1-2 IBM Diamondback internal components rear view in Figure 1-3

Number	Component
1	Library frame
2	Rear door
3	Rear door safety switch
4	IBM LTO 9 tape drives
5	Library controller (electronics tray)
6	10 cartridge service magazine
7	Rear view of ultra deep cartridge slots

Figure 1-4 shows a close view of the electronics tray, with both network ports, technician port and the TSSC port visible to the left and the dual redundant power supplies on the right.



Figure 1-4 IBM Diamondback electronics tray close-up

Figure 1-5 and Figure 1-6 show the front and rear views of the robotic accessor assembly, and Figure 1-7 on page 17 shows a close up view of the door sensor card.




Figure 1-5 A close up front view of the robotic accessor assembly, including the tape grippers and bar code scanner.




Figure 1-6 A close up view of the rear of the robotic accessor assembly.




Figure 1-7 A close up view of the door sensor card

1.3 Feature codes for IBM Diamondback

You can order feature codes (FCs) to install components in a library frame or to enhance the capacity or capabilities of the library.

A client-setup unit, which is referred to as a CSU in some product documentation, is a feature that 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 a charge. However, detailed installation instructions are included when you order and receive these features.

The FCs that are available for the IBM Diamondback are listed in Appendix A, “Feature codes for IBM Diamondback” on page 129

1.3.1 External IBM Tape System Service Console (TSSC)

Diamondback supports connecting to an external TSSC. A single TSSC can provide Call Home capability for several Diamondback, TS7700, TS4500, and TS3500 devices at the same site as shown in Figure 1-8 on page 18.




Figure 1-8 Using an external TSSC for Diamondback and other IBM tape libraries

An integrated management console (IMC) inside another device, for example, a TS4500, can also be used as an external TSSC. Feature code 2704 provides a 26-port switch. A maximum of 24 additional devices can be connected. Only one TSSC/IMC can be configured for IBM Call Home for a device. For more information about the TSSC, see 5.8, "Remote support" on page 151.

1.3.2 Top of rack frame

The IBM Diamondback library allows for a top of rack frame to be added, which provides 5U of available space. FC 1553 provides a top rack, which conforms to the Open Compute Project Olympus rack specification, and feature code 1154 provides a top rack in conformance with industry standard 19in racks.

When it is included, the top rack is factory-installed, but it can be easily removed during library installation if site doors are not tall enough to allow the library to easily pass through.

1.4 Host platforms and device drivers

The IBM Diamondback tape library is supported by a wide variety of servers, operating systems, and adapters. There are several ways to determine the servers and software that are supported.

To find information about the list of compatible components, several options are available:

- ▶ For a current list of compatible software, operating systems, and servers for LTO tape drives, see [Independent Software Vendor \(ISV\) matrix for IBM 3592 tape drives and LTO](#).
- ▶ For information about the IBM Diamondback tape library in a storage area network (SAN) configuration and for details on supported operating systems, servers, switches, and adapters, see [IBM System Storage Interoperation Center \(SSIC\)](#).
- ▶ To download any required tape device drivers or library firmware, see [Fix Central](#).
- ▶ For additional questions contact your IBM Sales rep directly, use the “Lets talk” web chat feature on [IBM.com](#) or call 1800-IBM-SERV.

1.5 IBM Diamondback specifications

This section contains physical dimensions and the environment and power requirements for the IBM Diamondback Tape Library.

1.5.1 Physical specifications

The following list includes the physical specifications of the tape library with and without the top rack:

- ▶ Width (with covers): 600 mm (23.62 in.)
- ▶ Depth (with doors): 1225 mm (48.22 in.)
- ▶ Height: 42U or 2025 mm (79.72 in.)
- ▶ Height (with top rack): 48U or 2295 mm (90.35 in.)
- ▶ Empty weight (without top rack): 413 kg (908 lb)
- ▶ Empty weight (with top rack): 438 kg (964 lb)
- ▶ Fully loaded weight (without top rack): 785 kg (1,727 lb)
- ▶ Fully loaded weight (with top rack): 879 kg (1,933 lb)

1.5.2 Floor space requirements

Table 1-3 provides the floor load ratings and required weight-distribution areas.

Table 1-3 Floor load ratings and required weight-distribution areas Operating Environment

Configuration	Maximum weight of configuration ^a	Floor load rating kg/m (lb/sq ft) ^b	Front weight distribution area mm (in.) ^{c d}	Rear weight distribution area mm (in.) ^{c d}
4884-L9A (without Top-Rack)	785 kg (1727 lbs)	610 (125)	460 (18.1)	460 (18.1)
		488 (100)	730 (28.7)	730 (28.7)
		439 (90)	880 (34.6)	880 (34.6)
		342 (70)	1300 (51.2)	1300 (51.2)

Configuration	Maximum weight of configuration ^a	Floor load rating kg/m (lb/sq ft) ^b	Front weight distribution area mm (in.) ^{c d}	Rear weight distribution area mm (in.) ^{c d}
4884-L9A (with Top-Rack)	879 kg (1933 lbs) ⁽²⁾	610 (125)	590 (23.2)	590 (23.2)
		488 (100)	890 (35.0)	890 (35.0)
		439 (90)	1060 (41.6)	1060 (41.6)
		342 (70)	1530 (60.2)	1530 (60.2)

- a. Maximum weight includes the highest capacity of tape drives, tape cartridges, and other feature codes that are possible to be installed in the tape library.
- b. This weight assumes that the Top-Rack is fully loaded with 68.2 kg (150 lbs.) of additional equipment.
- c. No side clearance is required for the tape library. Front and rear service clearance of 762 mm (30 in) is always required
- d. Weight distribution areas cannot overlap with weight distribution areas of nearby equipment. Service clearances that are described in note 3 are allowed to overlap with service clearances of adjacent equipment

The details that are provided in Table 1-4 assumes that the product is removed from the shipping container and installed. Also, the listed equipment environment specifications are for values that are measured at the cool air inlet:

Table 1-4 IBM Diamondback operating environment

	Allowable ^a	Recommended ^b	Maximum rate of change
Dry bulb temperature	16°C–32°C (60°F–90°F)	16°C– 25°C (60°F– 77°F)	5°C (9°F) per hour
Humidity range, non condensing	20% to 80% RH \oplus	20% to 50% RH	5% RH per hour No condensation ^c
Maximum wet-bulb temperature ^d	26°C 70°F		
Maximum dew point temperature ^e	22°C 72°F		
Maximum Elevation	3050m (10,000 ft)		

- a. Derate maximum dry-bulb temperature 1°C per 300 m above 900 m (1.8°F per 1,000 feet above 3,000 feet)
- b. Derate maximum recommended dry-bulb temperature 1°C per 300 m above 1,800 m (1.8°F per 1,000 feet above 6,000 feet)
- c. For 3592 media, changes of up to 40% RH in 5 minutes are allowed as long as the 20%–80% absolute limits are not exceeded.
- d. Applies to LTO drive generations 1–8 and to legacy 3592 drives (TS1155 and prior generations)
- e. Applies to TS1160 and LTO 9 drives

1.5.3 Power and cooling

Table 1-5 BM Diamondback power consumption

	Power Consumption (watts)			Cooling Requirements
	Off ^a	Idle	Maximum continuous	BTU per hour Maximum continuous ^{b(2)}
Library without drives or PDU (frame only)	11	95	130	443
Each FC 1852 PDU (TOR)	17	17	17	58
Each FC 1852 PDU (in frame)	17	17	17	58
Each LTO 9 Drive ^{c(3)}	0	18	37	126

- a. 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.
- b. To calculate the total cooling required by the library in Btu/hr, multiply the total power in watts by 3.41. To convert Btu per hr to kBtu per hour, divide your result by 1000.
- c. 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 includes 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.



IBM Diamondback Linear Tape-Open

The IBM Diamondback tape library supports Linear Tape-Open (LTO) 9 cartridges. The following tape drive models are supported:

- ▶ LTO 9 Full High (FH) Fiber Channel Drive

The IBM LTO 9 tape drive supports the LTO Generation 9 media specification of more than a 50% increase of the effective storage capacity to 45 TB with 2.5:1 compression (18 TB native capacity) per tape cartridge. The capacity of the earlier LTO 8 effective is a maximum of 30 TB with 2.5:1 compression (12 TB native capacity) per tape cartridge. IBM LTO 9 tape drives can read and write LTO 8 data cartridges. The IBM LTO 9 has an 8 Gbps Fiber Channel interface for connection to a wide spectrum of open system servers.

- ▶ LTO 9 FH (Full High) SAS

The IBM LTO 9 Tape Drive supports the LTO Generation 9 media specification over 50% increase of the effective storage capacity to 45 TB with 2.5:1 compression (18 TB native capacity). The effective capacity of the LTO 8 tape drive is a maximum of 30 TB with 2.5:1 compression (12 TB native capacity) per tape cartridge. IBM LTO 9 tape drives can read and write LTO 8 data cartridges. The IBM LTO 9 has a 12 Gbps SAS interface for connection to a wide spectrum of open system servers.

This chapter includes the following topics:

- ▶ 2.1, “IBM LTO tape drives for the Diamondback library” on page 34
- ▶ 2.2, “IBM LTO 9 tape drives” on page 46

2.1 IBM LTO tape drives for the Diamondback library

The three sponsoring companies of the LTO format protect the investments of clients by providing a roadmap to LTO generation 12. Figure 2-10 illustrates the roadmap for native cartridge capacity Figure 2-1. The three LTO sponsoring companies also established an infrastructure to enable compatibility between products. At the time of this writing, five generations are available.

Important Note: IBM, Hewlett-Packard, and Quantum reserve the right to change the information in this roadmap (Figure 2-1) without notice.

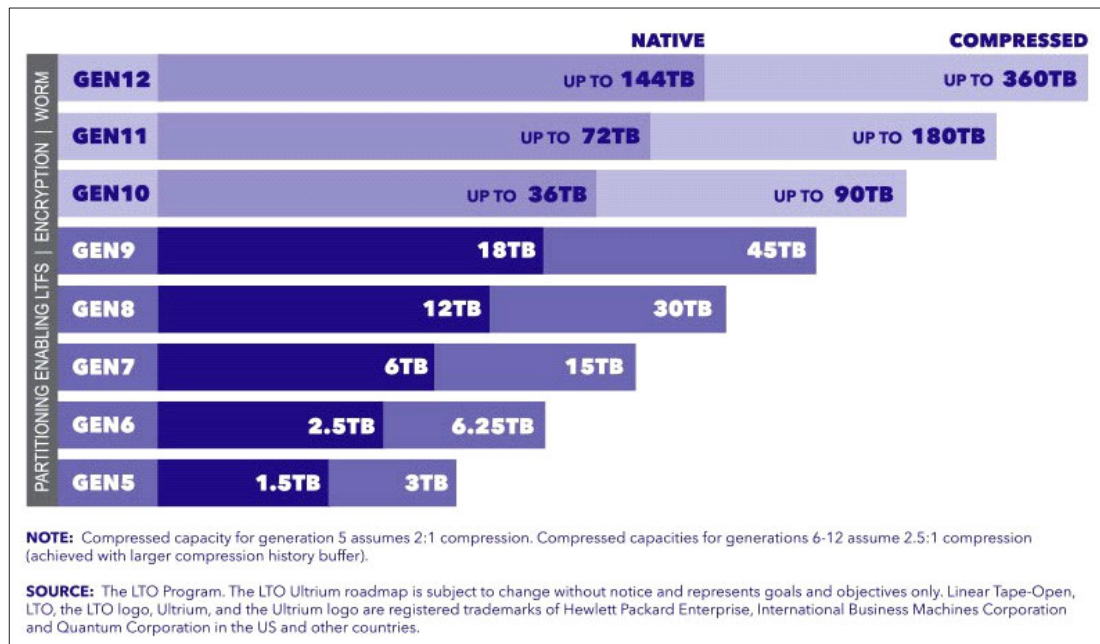


Figure 2-1 LTO Roadmap

Note: For more information, see [What is LTO Technology?](#)

The three technology provider companies, IBM, HP, and Quantum, all made significant contributions of time and expertise to define the LTO format specifications. All three companies have deep knowledge of clients' needs. They provided expert knowledge and engineering skills in the critical areas of magnetic recording technology, mechanism design, media materials, and cartridge design. This cooperative process resulted in the creation of stronger LTO format definitions.

2.1.1 LTO core technology

MultiFibre Channel 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 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 9 format, the 9860 tracks are split into four data bands of 2465 tracks each, and 32 read/write channels. The values for LTO 9 and the previous LTO generations are listed in Table 2-1.

Table 2-1 Data tracks, density, and channels for LTO generations

LTO generation	Data tracks	Tracks/band	Linear density	Read/write channels
LTO 1	384	96	124 Kbps	8
LTO 2	512	128	188 Kbps	8
LTO 3	704	176	250 Kbps	16
LTO 4	896	224	328 Kbps	16
LTO 5	1,280	320	368 Kbps	16
LTO 6	2,176	544	385 Kbps	16
LTO 7	3,584	896	485 Kbps	32
LTO 8	6,656	1,664	524 Kbps	32
LTO 9	9,860	2,465	545 Kbps	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. The center is the most physically stable area on the tape. Data is also 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 9 format are concurrently read or written. See the Read/write channels column in Table 2-1. 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 are high (Table 2-1) 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 in the 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 provides storage and retrieval. The following list describes the nonvolatile memory capacity of each drive type:

- ▶ 32,640 bytes for LTO 9
- ▶ 16,320 bytes for LTO 8, 7 and 6
- ▶ 8192 bytes for LTO 5 and 4

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.

2.1.2 Data compression

The LTO Consortium created a superior data compression technique that is known as LTO Data Compression (LTO-DC). Although it is 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, also known as 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 2-2 on page 37 shows the LTO-DC data compression technique that uses the two schemes.

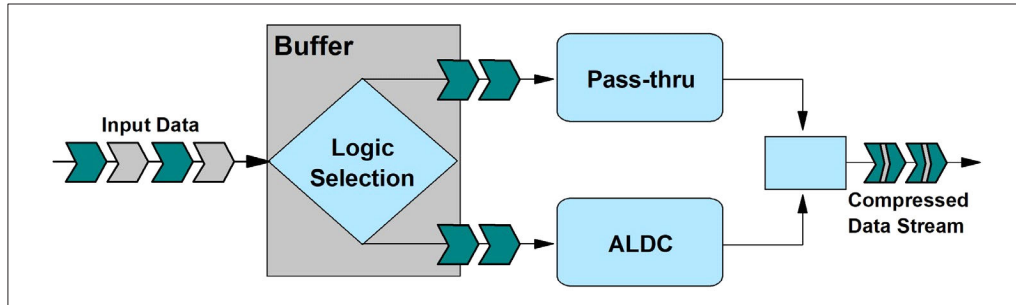


Figure 2-2 LTO-DC block diagram

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 [ECMA-321 Streaming Lossless Data Compression Algorithm](#).

Note: LTO uses the SLDC technique for compression.

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 LTO drive can read and decompress the LTO-DC streams of the other vendors' LTO drives.

2.1.3 LTO tape media

The IBM LTO cartridges are distinguished by color:

- ▶ LTO 6 cartridge is black.
- ▶ LTO 7 cartridge is purple.
- ▶ LTO 8 cartridge is burgundy.
- ▶ LTO 9 cartridge is green.

The IBM write-once, read-many (WORM) cartridges are two-tone cartridges with a platinum bottom:

- ▶ LTO 6 cartridge is black with a platinum bottom.
- ▶ LTO 7 cartridge is purple with a platinum bottom.
- ▶ LTO 8 cartridge is burgundy with a platinum bottom.
- ▶ LTO 9 cartridge is green with a platinum bottom.

Data cartridge

The LTO 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 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.

LTO drive technology is intended to meet the needs of the enterprise on a roadmap, or migration path, that extends well into the future. The LTO tape format established a new benchmark for large volume backup and archive options.



Figure 2-3 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 has a platinum bottom. All generations of cartridges contain half-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) and 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) and 30 TB (27.93 TiB) at 2.5:1 compression.

WORM tape format

Beginning with LTO format generation 3, WORM functionality provides non-erasable, non-rewritable operation with tape media. WORM is for long-term, tamper-resistant record retention. Starting with LTO 4 drives, all LTO drives of the same generation provide the same WORM capability.

The format specification for WORM for LTO 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 Diamondback 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 recommend for use on the Diamondback library. They are most often the cause of reading media label issues.

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 Lx, where x equals 1, 2, 3, 4, 5, 6, 7, 8, 9, T, U, V, W, X, or Y, which identifies the cartridge as an LTO cartridge. M8 media type denotes an LTO 7 cartridge that is 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 2-4 show a sample bar code label for an LTO 9 tape cartridge.



Figure 2-4 Example of an LTO 9 tape cartridge label and its components

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.

For more information about bar code label specifications, see [IBM Support: IBM LTO Ultrium Cartridge Label Specification](#).

Cleaning cartridges

The IBM LTO Cleaning Cartridge, which is also known as the universal cleaning cartridge, and the LTO 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 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 [IBM TotalStorage LTO Ultrium Tape Drive SCSI Reference \(LTO-5 through LTO-9\)](#).

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 memory of 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, which is twice 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 is also 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. Previously, information was written on the tape medium in a portion of the tape before the beginning-of-tape marker that users were not allowed to access, 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 Diamondback tape library are listed in Table 2-2.

Table 2-2 LTO 9 tape drive cartridge compatibility

Cartridge Type	Drive Type
	LTO 9
18 TB LTO 9 Data Cartridge (xxxxxxL9)	R/W
18 TB LTO 9 WORM Cartridge (xxxxxxLZ)	R/W
12 TB LTO 8 Data Cartridge (xxxxxxL8)	R/W
12 TB LTO 8 WORM Cartridge (xxxxxxLY)	R/W
9 TB LTO M8 Data Cartridge (xxxxxxxM8)	Not compatible
6 TB LTO 7 Data Cartridge (xxxxxxL7)	Not compatible
6 TB LTO 7 WORM Cartridge (xxxxxxLX)	Not compatible
2.5 TB LTO 6 Data Cartridge (xxxxxxL6)	Not compatible
2.5 TB LTO 6 WORM Cartridge (xxxxxxLW)	Not compatible
1.5 TB LTO 5 Data Cartridge (xxxxxxL5)	Not compatible
1.5 TB LTO 5 WORM Cartridge (xxxxxxLV)	Not compatible
800 GB LTO 4 Data Cartridge (xxxxxxL4)	Not compatible
800 GB LTO 4 WORM Cartridge (xxxxxxLU)	Not compatible
400 GB LTO 3 Data Cartridge (xxxxxxL3)	Not compatible
400 GB LTO 3 WORM Cartridge (xxxxxxLT)	Not compatible

Note: The drive rejects any command to load unsupported media and returns a sense key of 3, with an extra sense code/additional sense code qualifier of 30/01.

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. If you use M8 media with earlier firmware levels, then the M8 media is configured as an L8 volume. This can result in a permanent reduction of the capacity of the M8 media from 9 TB to 6 TB.

2.1.4 LTO 9 media optimization

Media optimization is a new feature for the LTO 9 tape drive with L9 and LZ media.

LTO 9 media requires greater precision because data can be written on an increased number of tracks. Media optimization creates a referenced calibration for each cartridge, which is used to align the tape drive to optimize data placement across these tracks. LTO 9 media optimization also enhances LTO tape long-term media durability.

Higher track densities, which are measured as tracks per inch (TPI), and longer tapes that use thinner polymer base substrates in LTO 9 media achieve double-digit areal densities and more than 21,000 TPI. A new tape calibration algorithm was introduced that ensures data durability, compatibility, and allows for complex interchange requirements over the specified environmental conditions. The length of calibration is mainly for media optimum characterization and varies based on these factors. This initial calibration also enables the tape drive's intelligent alignment to optimize data placement.

It is important to understand when media optimization occurs:

- ▶ Media optimization is performed on the first load of media that is labeled as L9 and LZ during initialization.
- ▶ It is a recommendation to perform the first load in the location of deployment, which should be in a stable environment that meets the recommended environmental specification.
- ▶ Media optimization is a one-time operation that can be completed on any drive in the environment. After media optimization, the media can be used across all tape drives without further optimization.

It is also important to consider the duration of media optimization:

- ▶ Media optimization averages 20 minutes per first load of a cartridge to a tape drive. Although most media optimizations complete within 30 minutes, some media optimizations can take a maximum of 2 hours.
- ▶ Interruption of the process is not recommended.
- ▶ A different mount does not necessarily improve the time to complete the one-time optimization.

Note: When using LTO 8 media, there is no media optimization. The media can be optimized by using the specific settings of a tool such as ITDT. The result is the loss of all existing data.

Some applications might require an update if the optimization time is too long for the current version of the application. Contact your software application provider for more details. Customized software that is not provided as a standard market product might require modification to ensure that the software can handle the extended first mount and optimization time.

2.1.5 Physical attachment

The IBM LTO 9 drive 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 or 12 Gbps SAS connectivity.

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. For more information, see [IBM System Storage Operation Center](#).

The Diamondback LTO 9 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 tape drive is shown in Figure 2-5.




Figure 2-5 LTO 9 drive front panel

The back panel of the LTO 9 tape drive is shown in Figure 2-6 on page 44.

The components on the rear panel and the normal status of the LEDs are listed in Table 2-3 on page 44:

- ▶ Number 1 in the picture represents the type of LTO.
- ▶ Number 8 is the fiber channel port number: port 0.
- ▶ Number 9 is the Fibre Channel port number: port 1




Figure 2-6 LTO 9 drive rear panel

Table 2-3 LED status for the LTO 9 drive.

LED	LED name	LED color	LED On	LED Flashing	LED Off
2	Drive Power	Green	+5 V dc or +12 V dc is present at the drive	During a rest, the LED might flash, which means that the drive is trying to unload a cartridge before processing the reset.	+5 V dc or +12 V dc is missing at drive.
3	Drive Status	Yellow	<ul style="list-style-type: none"> ▶ Cartridge status not indicated ▶ No message or error ▶ POR Diagnostic in progress 	<ul style="list-style-type: none"> ▶ Cartridge not present ▶ Message or error ▶ POR completed 	<ul style="list-style-type: none"> ▶ Cartridge not present ▶ No message or error ▶ POR diagnostic complete
		Green	<ul style="list-style-type: none"> ▶ Cartridge present ▶ No message or error ▶ POR diagnostic complete 	<ul style="list-style-type: none"> ▶ Cartridge present ▶ Message or error ▶ POR diagnostic complete 	
4	Library Communication	Green	At least one message received since POST	Library to drive communication activity	No communications since last POST of drive

LED	LED name	LED color	LED On	LED Flashing	LED Off
5	Information	Blue	There is a problem. More information is available on the management GUI error page.	N/A	N/A
6	Port 0 - Fibre Channel Activity	Yellow	N/A	Fibre Channel connections are good; not ready to transmit	<ul style="list-style-type: none"> ▶ No light detected through Fibre Channel. ▶ Fibre Channel cable is disconnected ▶ Port is not in use. ▶ Fibre wrap plug is installed on the port. ▶ Drive was not returned to operator mode.
7	Port 1- Fibre Channel Activity	Green	N/A	Fiber Channel connections are good; data is being transmitted.	

Figure 2-7 shows the LTO 9 SAS tape drive. Except for the physical interface, the color of the back of the LTO 9 SAS tape drive is gray.

The back of the LTO 9 Fibre Channel tape drive is black. The LED indicators are the same as the LTO 9 Fibre Channel drive.



Figure 2-7 Back of the LTO 9 SAS drive

2.1.6 Performance specifications

Performance specifications of the Diamondback library LTO 9 drives are listed in Table 2-4. Only LTO 9 and LTO 8 media are supported.

Table 2-4 Performance specifications for LTO 9 tape drives

Performance Parameter	Tape Drive	
	LTO 9 (SAS)	LTO 9 (FC)
Sustained data rate (native)*	400 MBps (L9 media)	400 MBps (L9 media)
	360 MBps (L8 media)	360 MBps (L8 media)
Sustained data rate (L9 media compressed at 2.5:1 compression)	900 MBps (L9 media)	750 MBps (L9 media)
	750 MBps (L8 media)	750 MBps (L8 media)
Maximum sustained data rate (at maximum compression)	900 MBps (L9 media) 500 MBps (L8 media)	700 MBps (L9 media) 700 MBps L8 media)
Burst data rate for Fibre Channel drives	1200 MBps	800 MBps
Time to load, thread, and initialize a cartridge	17 seconds	17 seconds
Time to unload cartridge from load point	24 seconds	24 seconds
Average space record time from load point	45 seconds (L9 media)	45 seconds (L9 media)

* All sustained data rates depend on the entire data path (from data source to host system to tape drive).

2.2 IBM LTO 9 tape drives

As the time of writing, LTO 9 is the latest LTO generation, which was released in September 2021. The IBM LTO 9 offerings represent significant improvements in capacity, performance, and reliability over the previous generation, LTO 8, but still protects the client's investment in the previous technology.

The LTO 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 LTO 9 tape drive provides the following improvements over the LTO 8 models:

- ▶ The LTO 9 tape drive increases the native transfer speed to 400 MBps compared with 360 MBps for the LTO 8 tape drive.

- ▶ The LTO 9 FC tape drive compressed speed of 700 MBps for Fibre Channel interface and LTO 9 SAS tape drive compressed speed to 900 MBps for SAS interface.
- ▶ The LTO 9 tape drive increases the potential capacity of a cartridge in LTO 8 format.

2.2.1 LTO 9 tape drive features

The LTO 9 tape drive includes the following features:

- ▶ Maximum native data transfer rate of 400 MBps
- ▶ Maximum compressed data transfer rate of 700 MBps for model F9C/S9C
- ▶ LTO 9 data and WORM tape cartridge native physical capacity of up to 18 TB
- ▶ 8 Gbps multi-mode Fibre Channel attachment option for model F8C
- ▶ 8 Gbps single-mode Fibre Channel attachment option for model F8S
- ▶ 12 Gbps dual-mode SAS attachment option for model S9C
- ▶ 1024 MB internal data buffer
- ▶ Application-managed encryption support
- ▶ Library-managed encryption support
- ▶ 32 KB cartridge memory with LTO 9 media
- ▶ Open Recommended Access Order (oRAO) to LTO 9 technology
- ▶ IBM Storage Archive Single Drive Edition (SDE), Library Edition (LE), and Enterprise Edition (EE) support
- ▶ Provides partitioning support, which, with IBM Storage Archive, provides users with file-level access to tape data

2.2.2 Compatibility

LTO 9 drives can read and write on LTO 8. It cannot read earlier generations of LTO drives. LTO 9 drives can also read and write on LTO 8 cartridges at 12 TB capacity.

The drive also writes to tapes that can be read by other licensed LTO 9 drives. In addition to the use of the IBM LTO 8 data cartridge (with up to 12 TB capacity), the drive offers read/write capability for certified LTO 8 tape cartridges.

Note: The LTO 9 tape drive can read or write to LTO 9 and LTO 8 format cartridges only.

The native data transfer rate when a data cartridge of another generation is processed is listed in Table 2-5.

Table 2-5 Native data transfer rate with various media

	18 TB (LTO 9)	12 TB (LTO 8)
Supported methods of operation	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)

2.2.3 Connectivity

The LTO 9 tape drive family offers high capacity, performance, and technology for the midrange open systems environment. LTO 9 model F9C is a dual-ported tape drive that facilitates 8 Gbps Fiber 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 [IBM System Storage Interoperations Center \(SSIC\)](#).

2.2.4 Performance

The LTO 9 tape drive uses 8960 data tracks to read and write to LTO 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 LTO 9 SAS and FC tape drive models 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 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 feature was introduced to IBM LTO 9 tape drive technology, Open Recommended Access Order (oRAO). This technology reduces tape data access times in LTO Gen 9 tape technologies by a maximum of 73%.

oRAO enables tape control applications to accelerate the retrieval of a specific number of portions of a file, or file chunks, from a single tape by reducing the seek time between those file chunks. (Figure 2-8).

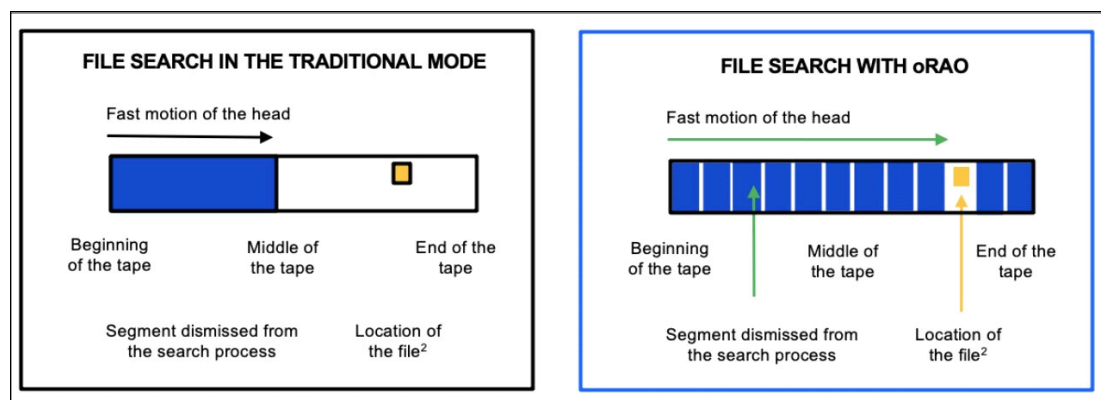


Figure 2-8 IBM LTO 9 Open Recommended Access Order vs previous method for data access

This feature was developed based on IBM file access acceleration technology. It can improve cyber resilience by optimizing the access times to recovery data and reducing read times. 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 retrieval by using the IBM LTO 9 tape drive oRAO technology when working with larger volumes of data (see Figure 2-8 on page 48).

Note: IBM LTO 9 oRAO is available on LTO 9 Full High tape drives only.

2.2.5 Dynamic speed matching

The LTO 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 9 drives are listed in Table 2-6.

Table 2-6 LTO 9 drive speed matching data rates

Performance parameters	Generation 9 media	Generation 8 media
Speed matching data rates in MBps	408	365.0
	385	341.0
	366	318.0
	347	306.4
	325	273.0
	305	249.5
	284	226.0
	263	203.0
	244	180.0
	223	157.5
	203	135.0
	177	112.0

2.2.6 Encryption

The LTO 9 tape drive family is encryption-capable and supports application-managed encryption (AME) at no charge on Fibre Channel or SAS tape drives.

2.2.7 IBM Storage Archive software application

IBM LTO 9 tape drives are compatible with the IBM Storage Archive software application and the underlying linear tape file system (LTFS). LTFS uses the LTO media partitioning functionality, which is present on LTO 9, LTO 8, LTO 7, LTO 6, and LTO 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 Storage Archive supports the IBM LTO 9, 8, 7, 6, and 5, and IBM TS1170, TS1160, TS1155, TS1150, and TS1140 tape drives. IBM Storage Archive supports the LTO 7 initialized LTO 8 Type M cartridge.

Tape as a storage medium offers many benefits. It is reliable, portable, low-cost, low-power, and high-capacity. However, tape is typically not as easily accessible 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 is comparable to using a USB flash drive. With LTFS file manager, reading data on a tape cartridge can be done by 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 Storage Archive supports IBM tape automation in addition to the single drive edition. By using IBM Storage 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.

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.

2.2.8 Surface Control Guiding Mechanism

The *Surface Control Guiding Mechanism* is designed to guide the tape along the tape path in the LTO 6 tape and later drives. This method uses the surface of the tape rather than the edges to control tape motion.

LTO 6 and later 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.

2.2.9 Statistical Analysis and Reporting System (SARS)

The LTO 9 tape drives communicate with the Diamondback 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.

2.2.10 Media partitioning

Media partitioning is now available with the LTO 6 and later tape drive. Media partitioning allows for faster data access by splitting the cartridge into two media partitions. WORM media is not partitionable.

2.2.11 IBM Storage Archive software application

IBM LTO 6 and later tape drives are compatible with the IBM Storage Archive software application. IBM Storage Archive and the underlying Linear Tape File System (LTFS) use the LTO 6 and 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 LTO 9, 8, 7, 6, and 5, and IBM TS1170, 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 always easy to use. It has no standard format, and data often cannot be used without first copying the data to disk.

IBM Storage Archive supports IBM tape automation in addition to the single drive edition. By using IBM Storage 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. Also, it provides for caching of tape indexes and for searching, querying, and displaying the contents of tapes within an IBM tape library without the requirement to mount tape cartridges.



IBM Diamondback Tape Library Supported Software

This chapter highlights some of the software solutions that are commonly used with the IBM Diamondback Tape Library.

- ▶ 3.1, “IBM Storage Archive” on page 34
- ▶ 3.2, “Linear Tape File System” on page 35

3.1 IBM Storage Archive

IBM Storage Archive (formally IBM Spectrum Archive) is a member of the IBM Storage family. IBM Storage Archive enables direct, intuitive, and graphical access to data that is stored in IBM tape drives and libraries. IBM Storage Archive uses the Linear Tape File System (LTFS) format standard for reading, writing, and exchanging descriptive metadata on formatted tape cartridges.

IBM Storage Archive eliminates the need for more tape management and software to access data.

IBM Storage Archive offers three software solutions for managing your digital files with the LTFS format:

1. Single Drive Edition (SDE)
2. Library Edition (LE)
3. Enterprise Edition (EE)

By using IBM Storage Archive Enterprise Edition and IBM Storage Scale, a tape tier can be added 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.

3.1.1 Key capabilities

IBM Storage 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 Storage Scale

3.1.2 Benefits

IBM Storage Archive takes advantage of the low cost of tape storage. IBM Storage Archive includes the following benefits:

- ▶ Enable disk-like 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
- ▶ Drag and drop interface for selecting tapes

3.2 Linear Tape File System

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 Storage 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 Storage Archive SDE is for single TS1170, TS1160, TS1155, TS1150, TS1140, LTO 9, LTO 8, LTO 7, LTO 6, or LTO 5 tape drives. IBM Storage Archive SDE can be downloaded from [IBM Fix Central](#).

LTFS is the first file system that works with LTO generation 9, 8, 7, 6, 5, TS1170, TS1160, TS1155, TS1150, and TS1140 tape technology. LTFS sets a new standard for ease of use and portability for open systems tape storage. With LTFS, accessing data that is stored on an IBM tape cartridge is comparable to 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 (OS) 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 because of the addition of file tagging.
- ▶ Files can be moved to and from LTFS tape by using drag-and-drop.
- ▶ 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, which is affected by 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 systems, which are fully compatible with LTFS.

Important: LTO cartridges from earlier LTO generations (that is, LTO 1 through LTO 4) are not partitionable. Therefore, they cannot be used by LTFS. If LTO 4 cartridges are used in an LTO 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 5 drive.

Similarly, WORM cartridges of any generation cannot be used by LTFS because they cannot be partitioned.

The TS1170, TS1160, TS1155, TS1150, and TS1140 tape drives are also supported by IBM Storage Archive SDE, IBM Storage Archive LE, and IBM Storage 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.

When you use LTFS, you can simultaneously copy two or more files to an LTFS tape, but 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, TS1170, TS1160, TS1155, TS1150, or TS1140 tape drives and tape media

Together, these components operate on a file system on LTO media as though it is a disk file system for accessing tape files, including the tape directory tree structures. After it is mounted, the metadata of each cartridge 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.

3.2.1 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 IBM Storage Archive Single Library 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. Users can run any application that is designed for disk files against tape data without concern that the data is physically stored on tape. 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 Storage Archive SDE, see *IBM Spectrum Archive Single Drive Edition and Library Edition Installation and Configuration*, SG24-8090.

The TS1170, 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 TS1170, 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 TS1170, TS1160, TS1155, and TS1150 drives support both TS1140 style partitions.

The TS1170, TS1160, TS1155, and TS1150 support more partitioning enhancements over the TS1140, specifically, both wrap-wise partitioning and longitudinal partitioning methods.

Restriction: Longitudinal partitioning is only supported with media types before JE/JM.

The following characteristics apply in general to partitioning support on the TS1170 (3592 60F, 60S), TS1160 (3592 60E, 60F, and 60S), TS1155 (3592 55E and 55F), and TS1150 (3592 EH8):

- ▶ Partitioning is supported on media in the TS1170, 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 TS1170, 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 Tape Drive 3592 SCSI Reference](#)

The following types of partitioning are available:

- ▶ Wrap-wise partitioning (which is used on LTO 9, LTO 8, 7, 6, and 5, TS1170, TS1160, TS1155, TS1150, and TS1140).
- ▶ Longitudinal partitioning (a maximum of two partitions) can be used on TS1155, TS1150, and TS1140 only.

3.2.2 Wrap-wise partitioning

LTO 9, 8, 7, 6, and 5 support two wrap-wise partitions. The TS1170, TS1160, TS1155, TS1150, and TS1140 support four wrap-wise partitions.

Figure 3-1 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 3-1 Wrap-wise partitioning example

3.2.3 IBM Storage Archive Editions

As shown in Figure 3-2 on page 39, IBM Storage Archive is available in three editions that support small, medium, and enterprise businesses.

IBM Storage Archive Single Drive Edition

IBM Storage 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 Storage Archive Single Drive Edition file system works with LTO generation 9, 8, 7, 6, 5, TS1170, 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.

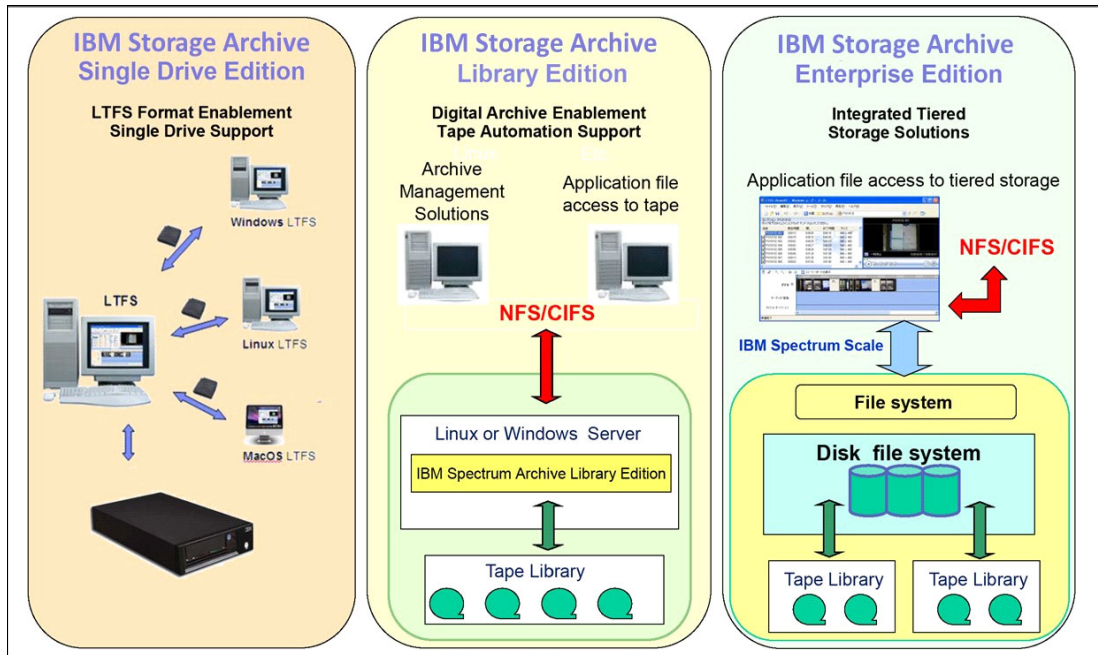


Figure 3-2 IBM Storage Archive SDE, LE, EE Editions

For more information about IBM Storage Archive SDE, see *IBM Linear Tape File System Installation and Configuration*, SG24-8090.

IBM Storage Archive Library Edition

IBM Storage Archive Library Edition (LE) extends the file manager capability of the IBM Storage Archive SDE. IBM Storage Archive LE is introduced with Version 2.0 of LTFS, enables 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 Storage 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 Storage Archive. IBM Storage 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
- ▶ Diamondback Tape Library

IBM TS1170, 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 Storage 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 Storage Archive LE supports both Linux and Windows.

IBM Storage Archive LE provides the following product features:

- ▶ Direct access and management of data on tape libraries with LTO 9, 8, 7, 6, 5 tape drives, and the TS1170, 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 Storage 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 open each of these folders to access the files that are stored on each tape. The IBM Storage 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 3-3 IBM Storage Archive LE view of tapes as folders

The following IBM tape libraries and tape autoloader support IBM Storage Archive LE:

- ▶ TS2900 Tape Autoloader
- ▶ TS3100 Tape Library
- ▶ TS3200 Tape Library
- ▶ TS3310 Tape Library
- ▶ TS4300 Tape Library
- ▶ TS3500 Tape Library
- ▶ TS4500 Tape Library
- ▶ Diamondback Tape Library

For more information about IBM Storage Archive LE, see *IBM Spectrum Archive Single Drive Edition and Library Edition Installation and Configuration*, SG24-8090.

IBM Storage Archive Enterprise Edition

IBM Storage Archive Enterprise Edition (EE) offers organizations a way to use cost-effective IBM tape drives and libraries within a tiered storage infrastructure. Instead of using disks for tier 2 and tier 3 data storage for data that is stored for long-term retention can use tape libraries to improve efficiency and reduce costs. In addition, IBM Storage Archive EE integrates with the scalability, manageability, and performance of IBM Storage Scale. IBM Storage Scale is an IBM enterprise file management platform that helps organizations move beyond adding storage to optimizing data management.

IBM Storage Archive Enterprise Edition offers the following benefits:

- ▶ Simplifies tape storage with the IBM Linear Tape File System (LTFS) format with the scalability, manageability, and performance of IBM Storage Scale
- ▶ Helps reduce IT expenses by replacing tier 2 and tier 3 disk storage with IBM tape libraries
- ▶ Expands archive capacity by adding and provisioning media without affecting the availability of data that is already in the pool
- ▶ Adds extensive capacity to IBM Storage Scale installations with lower media, less floor space, and lower power costs

IBM Storage Archive EE for the IBM Diamondback, TS4500, IBM TS3500, TS4300, and IBM TS3310 tape libraries provides seamless integration of IBM Storage Archive with IBM Storage Scale by creating an LTFS tape tier. You can run any application that is designed for disk files on tape by using IBM Storage Archive EE. IBM Storage 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 Storage Archive EE, you can enable the use of LTFS for the policy management of tape as a storage tier in an IBM Storage Scale environment and use tape as a critical tier in the storage environment.

IBM Storage Archive EE supports IBM Linear Tape-Open (LTO) 9, 8, 7, 6, and 5 tape drives, and IBM System Storage TS1170, TS1160, TS1155, TS1150, and TS1140 tape drives that are installed in Diamondback (LTO 9 only), TS4500 and TS3500 tape libraries or LTO 9, 8, 7, 6, and 5 tape drives that are installed in the TS3310 tape library or LTO 9, 8, 7, and 6 tape drives that are installed in the TS4300 Tape Library.

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

Figure 3-4 on page 42 shows the hybrid storage solution with the integration of the IBM Storage Archive EE with IBM Storage Scale.

The seamless integration offers 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

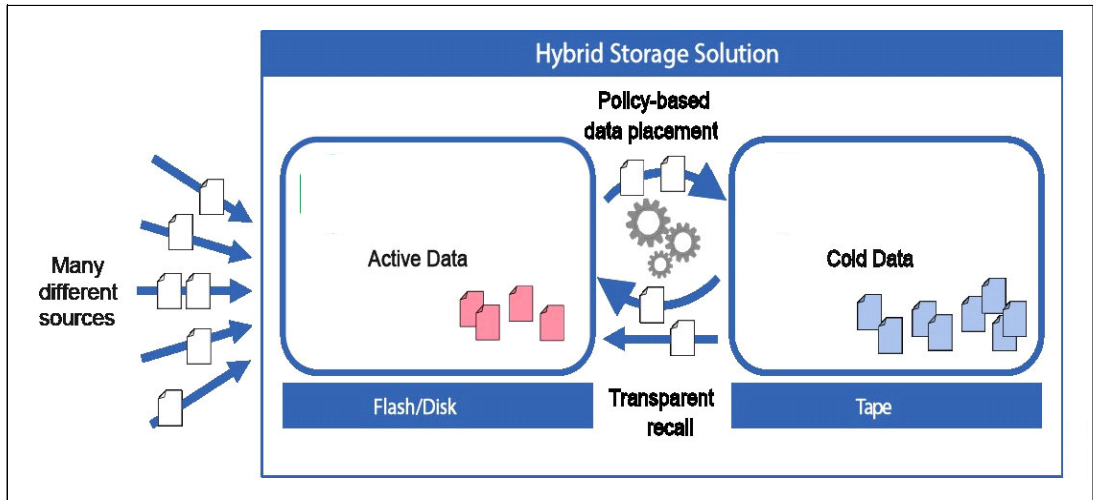


Figure 3-4 Integration of IBM Storage Scale and IBM Storage Archive Enterprise Edition

Note: For more information, see:

- ▶ [IBM Storage Archive Enterprise Edition \(EE\)](#)
- ▶ [IBM Storage Archive Library Edition \(LE\)](#)
- ▶ [IBM Storage Archive Single Drive Edition \(SDE\)](#)



Security and encryption

The IBM Diamondback tape library, as part of the IBM tape solution offerings, reflects years of commitment with technology and security. Things like Encryption policies, Air gaps, Worm cartridges are some of the different types of data protection that Diamondback can offer.

The encryption policies, methods, and software capabilities for the IBM Diamondback tape library are described in this chapter.

This chapter includes the following topics:

- ▶ 4.1, “Tape encryption overview” on page 50
- ▶ 4.2, “Encryption policy” on page 51

4.1 Tape encryption overview

The tape drives that are supported by the Diamondback 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.

4.1.1 Encryption-enabled tape drives

All of the tape drives that are supported by the Diamondback tape library are encryption-enabled.

4.1.2 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 Storage 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 Diamondback 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, [IBM Security Guardium Key Lifecycle Manager documentation](#).

4.2 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 4.2.2, “Encryption key management on the Diamondback” on page 52.

With the Diamondback tape library, the encryption policy is managed at the logical library level. The Logical Libraries page of the Diamondback management GUI is used to enable encryption for a logical library and modify the encryption method that is used. The Security page of the Diamondback management GUI is used to manage key servers and key labels.

Note: In the tape storage environment, the encryption function on tape drives that are on a desktop, are stand-alone, and are within libraries is configured and managed by the client. It is not configured and managed by the IBM service support representative (SSR).

4.2.1 Encryption methods

Encryption methods for the LTO 9 tape drives differ to some extent. The following sections include a brief description of these encryption methods. In these sections, the term *Key Manager (KM)* is used to refer to IBM Security Guardium Key Lifecycle Manager and any other key managers that might be used.

Symmetric key encryption

Encryption of data by using a symmetric key and algorithm is sometimes called *private key encryption* or *secret key*, which is different than 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. When you use IBM Security Guardium Key Lifecycle Manager, the same encryption key is used for encryption and decryption of the data. However, the 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.

4.2.2 Encryption key management on the Diamondback

The Diamondback tape library supports the use of IBM Security Guardium Key Lifecycle Manager for library-managed encryption (LME). For application-managed encryption, the application is responsible for managing its own keys.

The key manager operates on various 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. The key manager can serve numerous IBM encrypting tape drives, or encrypting disk drives, regardless of where those drives are installed. For example, the drives can be installed in tape library subsystems, which are connected to mainframe systems through various types of channel connections, or they can be installed in other computing systems.

The key manager uses a *keystore* to hold the certificates and keys or to hold 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 Diamondback tape library:

- ▶ Application-managed encryption (AME)
- ▶ 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 4-1 on page 53.

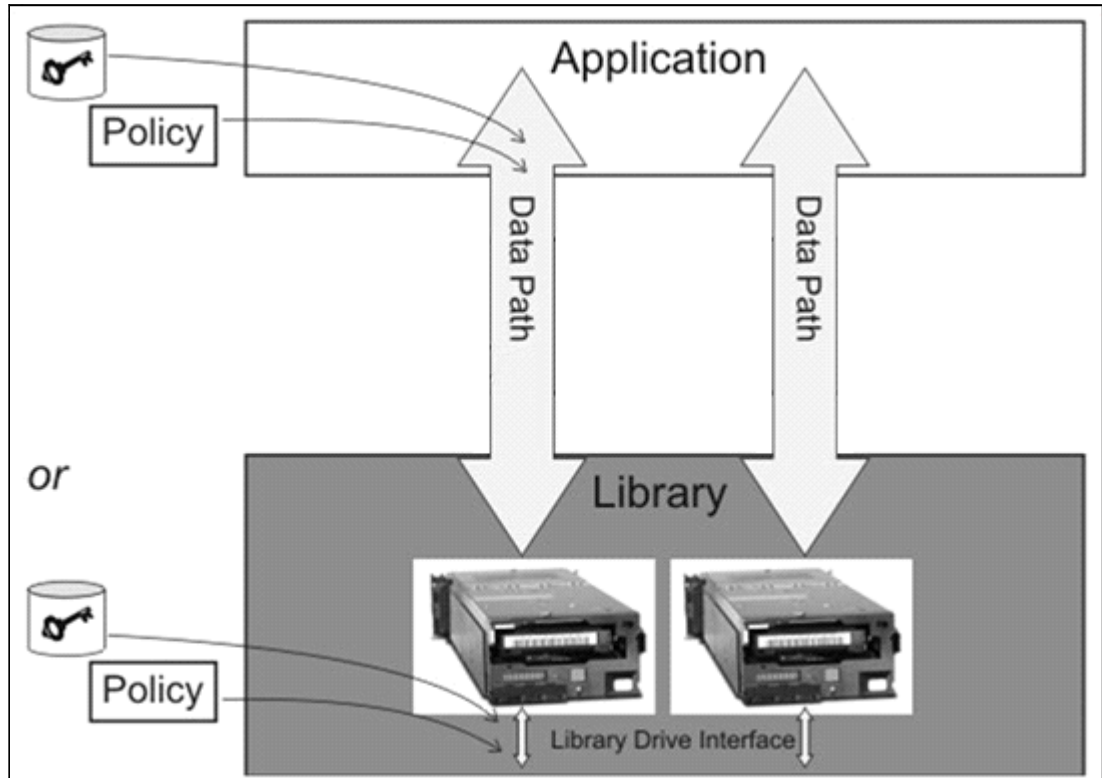


Figure 4-1 Possible locations for the encryption policy engine and key management

The application layer, for example, IBM Storage Protect, initiates the data transfer for tape storage.

The library layer is the Diamondback tape library, which contains an internal interface to each tape drive that is installed in the library.

4.2.3 Application-managed encryption

The AME method is best in operating environments that run an application that can generate and manage their own encryption policies and keys, such as IBM Storage 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 not noticeable 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:

1. The IBM encryption command set that was developed for the key manager
2. The T10 command set that was defined by the International Committee for Information Technology Standards (INCITS)

4.2.4 Library-managed encryption

Library-managed encryption (LME) is useful for encryption-enabled tape drives in an open-attached Diamondback 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 not noticed by the applications when it is used with certain applications, such as IBM Storage Protect.

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.

A maximum of four library-managed encryption (LME) key paths per logical library are supported on the Diamondback.

Note: 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

4.2.5 Prerequisites for using encryption on the IBM Diamondback tape library.

Certain hardware and software prerequisites must be met before you use encryption with the Diamondback tape library.

With the Diamondback 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 you use library-managed encryption (LME) or application-managed encryption (AME).

The following prerequisites must be met:

- ▶ Tape drives are enabled for encryption from the Logical Libraries page of the Diamondback management GUI.
- ▶ IBM Security Guardium Key Lifecycle Manager must be used as the key manager when you enable LME with LTO tape drives.

Adding an Encryption Key Server on the Diamondback

Use the Encryption Key Servers page to manage which key servers provide 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 and the Port that is to be used. If you use the DNS name, ensure that you added the DNS servers in the Network page.
4. Select the Security type: TLS1.2 or Proprietary.

If TLS1.2 is selected (see Figure 4-2), a Key server certificate must be added to the Diamondback by using the folder browser. The file can contain a chain of certificates. The file must also be in Base64 privacy enhanced mail (PEM) format. The maximum length is 255 characters.

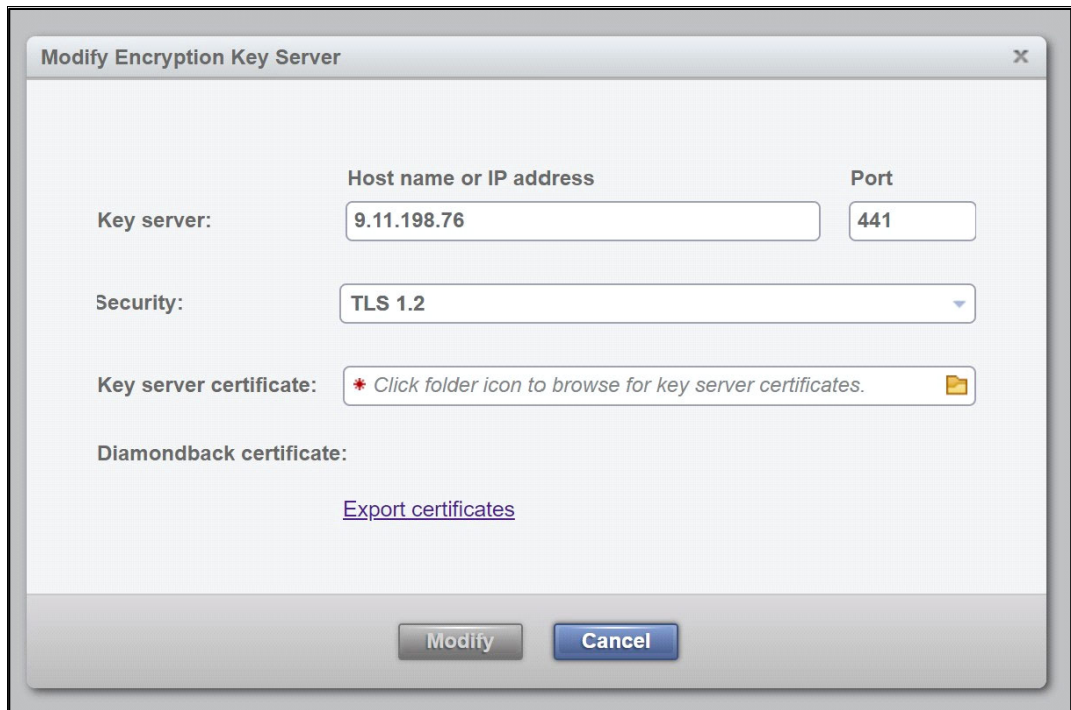


Figure 4-2 TLS 1.2 options

5. Click **Modify**.

Encryption methods on the Diamondback

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 5.3.1, “Creating a logical library” on page 81.

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 4-3 on page 56.

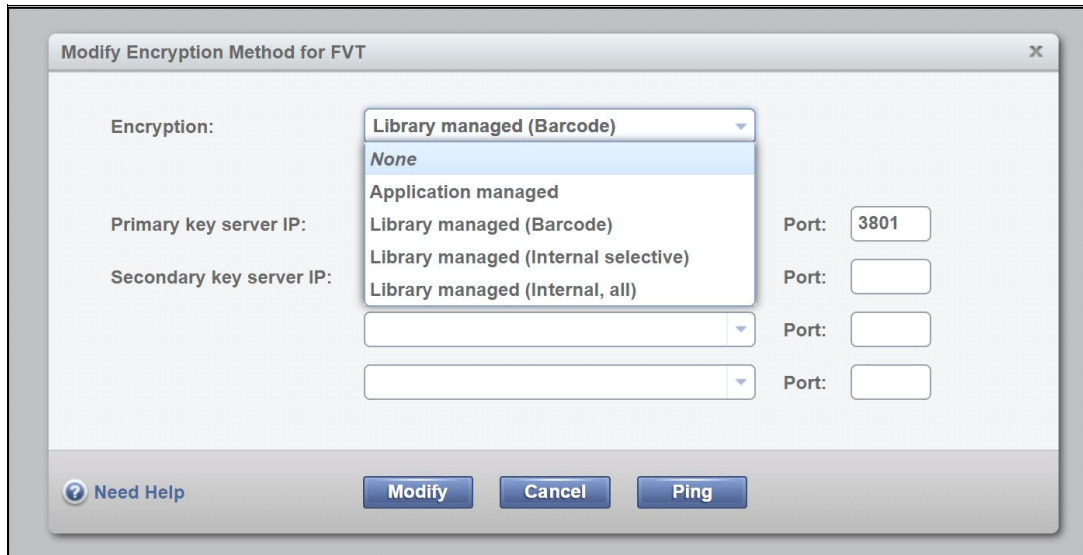


Figure 4-3 Modify Encryption Method window

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 Storage Protect, can manage encryption.

- ▶ 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, the 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.

4.3 Additional Security options

The use of physical tape, the Diamondback management software, and IBM LTO technologies allow for additional data security measures to be taken as needed. An overview of these additional security measures is outlined in the following section.

4.3.1 Air gap

During the COVID-19 pandemic, the number of ransomware attacks increased over 100%. Many companies and organizations that were not fully prepared had to choose between either losing data or paying large sums of money to bad actors, which did not always provide the decryption keys.

One major advantage of using tape as a backup is the physical separation from the organization's primary storage systems. The gap that is created between attackers and data, along with a good physical access policy, is one of the most effective solutions to recover from ransomware or any cyberattack in general.

Another advantage of tape is the ability to move the data during a disaster. Tape media can be moved while there is no power or network connections by transporting the data cartridge to a different data center to access and extract the information.

4.3.2 Safeguarded tape

Safeguarded tape is an orchestrated process, which creates an airgap, without physical media handling. By using the IBM Diamondback web-based GUI or REST interface, storage administrators can create logical libraries that are disconnected from all application access. Administrators can manually, or by using automation, connect to the Diamondback library and move selected media to the 'air gaped' logical library. The media is not physically moved, but is no longer visible or accessible to connected applications.

4.3.3 WORM tape

Write-once-read-many (WORM) cartridges are designed for applications such as archiving and data retention, and to prevent the alteration or deletion of user data. They are also suitable for applications that require an audit trail.

The supported LTO tape drives include the WORM feature, which is supported by the LTO WORM data cartridges. The following WORM media traits reduce tampering with data:

- ▶ The bottom of the WORM cartridge is molded in a color (gray) that is different from rewritable cartridges.
- ▶ A unique format is factory-written on each WORM cartridge.
- ▶ The WORM cartridge's memory, along with its unique format, protects the WORM character of the media.

Based on LTO technology, the format for the LTO WORM data cartridges provides the same capacities as the equivalent LTO data cartridges.



Diamondback management graphical user interface (GUI)

The Diamondback tape library provides advanced capabilities for integrated tape drive and media management, which is delivered within a single web-based GUI. Additionally, the Diamondback can be managed by accessing the library with a TSSC or integrated management console (IMC) which is part of another storage device. The Diamondback 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:

- ▶ 5.1, “Using the Diamondback management GUI” on page 54
- ▶ 5.2, “Monitoring” on page 58
- ▶ 5.3, “Library menu” on page 74
- ▶ 5.4, “Drives menu” on page 91
- ▶ 5.5, “Cartridges menu” on page 103
- ▶ 5.6, “Access menu” on page 118
- ▶ 5.7, “Settings” on page 127
- ▶ 5.8, “Remote support” on page 151
- ▶ 5.9, “IBM Net Promoter Score and Feedback” on page 162

5.1 Using the Diamondback management GUI

Before you connect to the management GUI, the Diamondback tape library must first be installed and configured either by the customer or by an IBM Customer Service Representative. Also, you must use one of the supported web browsers to access the Diamondback management GUI and enable cookies and JavaScript in your browser. Also, disable any browser software to block pop-up windows to ensure that all functions and features of the GUI work as expected.

5.1.1 Connecting to the management GUI

To connect to the management GUI, perform the following steps:

1. Open a web browser. The Diamondback GUI supports the latest 2 versions of the following web browsers:
 - Google Chrome
 - Mozilla Firefox or Mozilla Firefox ESR
 - Microsoft Edge
 - Safari
2. Enter the Ethernet IP address as the URL and press Enter.

Enter your username and password and click **Log in**. If this installation is new, log on to the Diamondback management GUI with the default username and password (User Name: admin and Password: admin) to continue with the setup. After logging in to the library with the default password, you will be prompted to create a new password for the account. The Diamondback management GUI login window opens as shown in Figure 5-1.

Important: Document the new administrator password. If additional administrator users are not defined, this password cannot be recovered unless you physically access the library hardware.

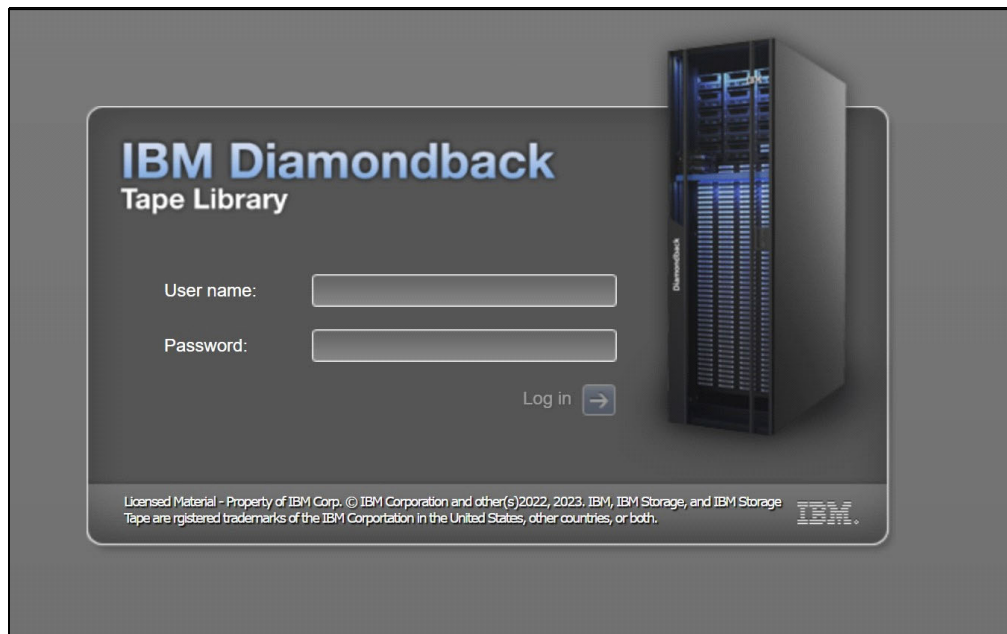


Figure 5-1 Diamondback Login Window

When performing a new installation of a Diamondback library, complete the following steps after the first login:

1. Set the local date and time. Select **Settings > Library > Date and time** (Figure 5-2).

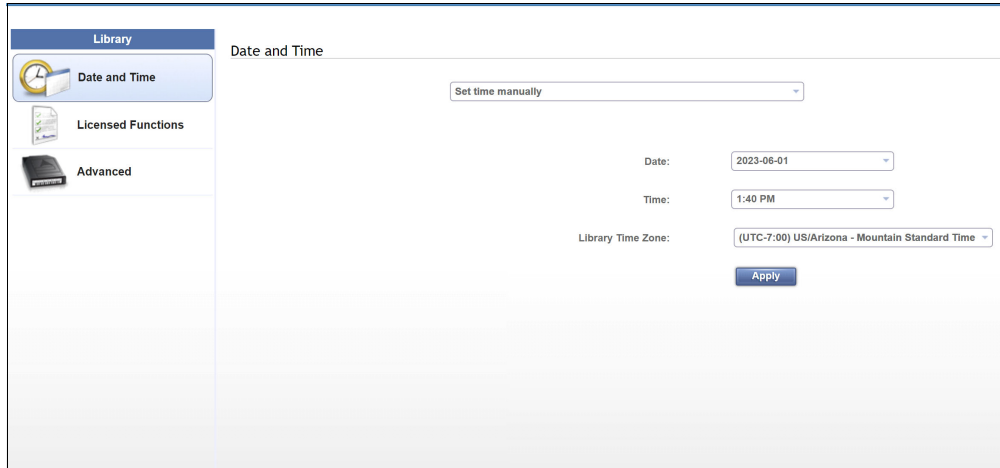


Figure 5-2 Set Date and Time

2. Calibrate the library by selecting **Actions > Calibrate Library** (Figure 5-3). It can take 11 minutes. You can Monitor the progress in the Tasks page.




Figure 5-3 Calibrate Library option

3. After Calibration is completed, select **Actions > Library Verify**.

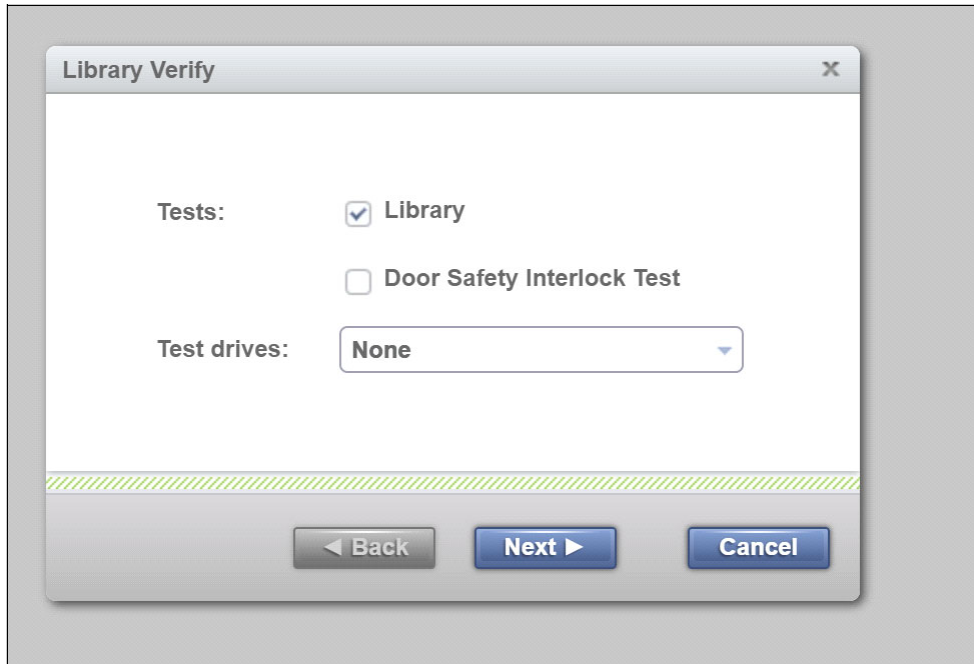


Figure 5-4 Verify Library command

After Library Verify is completed, the installation is finished.

5.1.2 System summary window

After you log in to the system, the system summary window is shown. (see: Figure 5-5).

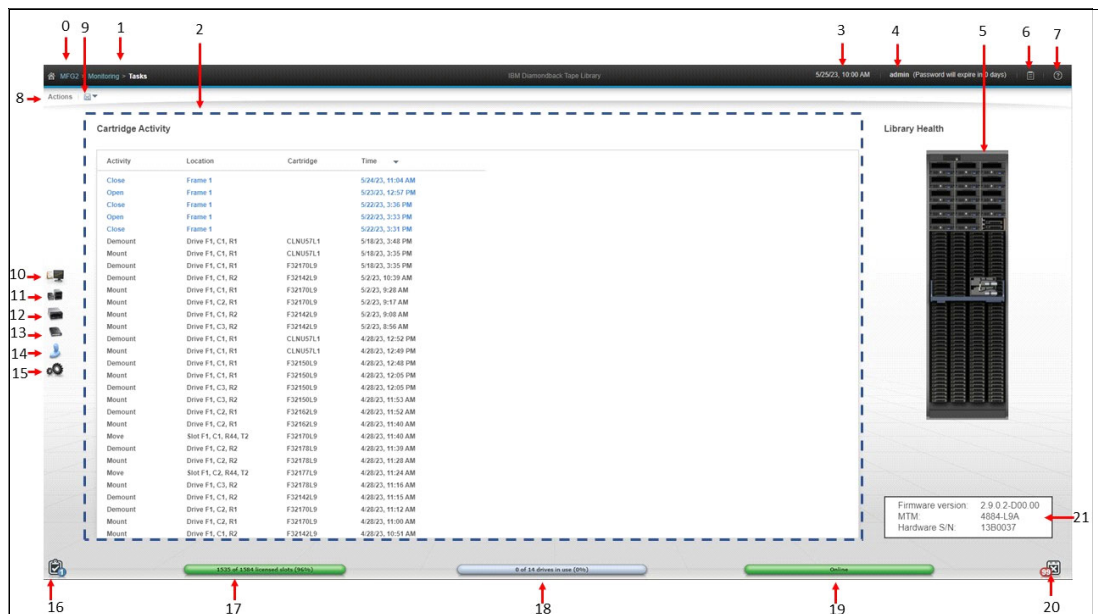


Figure 5-5 System summary window

For more information about the Management Interface (MI) and the available functions, use the question mark (?) icon or **Help**, which is number 6 in Figure 5-5.

Table 5-1 lists the fields that are shown in the window. The numbers in Figure 5-5 on page 56 refer to the numbers that are shown in Table 5-1.

Table 5-1 System display page components

Number	Field
0	IBM Diamondback tape library name.
1	Current menu tree position
2	The cartridge activity table.
3	Library Date and Time.
4	Logged-in username and role.
5	View of the physical library that is installed.
6	Suggested Tasks icon.
7	Help.
8	Library action pull-down menu.
9	Export Data menu.
10	Monitoring menu.
11	Library menu.
12	Drives menu.
13	Cartridges menu.
14	Access menu.
15	Settings menu.
16	Tasks icon. The Tasks icon displays when tasks are running.
17	The physical capacity pod shows the licensed capacity. This pod changes color to yellow or red, depending on whether the capacity exceeds the thresholds.
18	Drives in use status pod.
19 ^a	Library health status pod. The color of the health status pod indicates the current state of the library by severity.
20	Library events icon.
21	Library information box.

a. Status PODs. Pods are always 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.

5.2 Monitoring

The IBM Diamondback management GUI pages under the Monitoring icon can help you to monitor the library, events, and tasks, as shown in Figure 5-6.



Figure 5-6 Monitoring icon

5.2.1 Library

The Library page, as shown in Figure 5-7, displays the configuration (firmware version, Model, S/N) and the recent Activity of your IBM Diamondback Tape Library. Starting on this page is the best place to begin viewing the health and status of the library and library components.

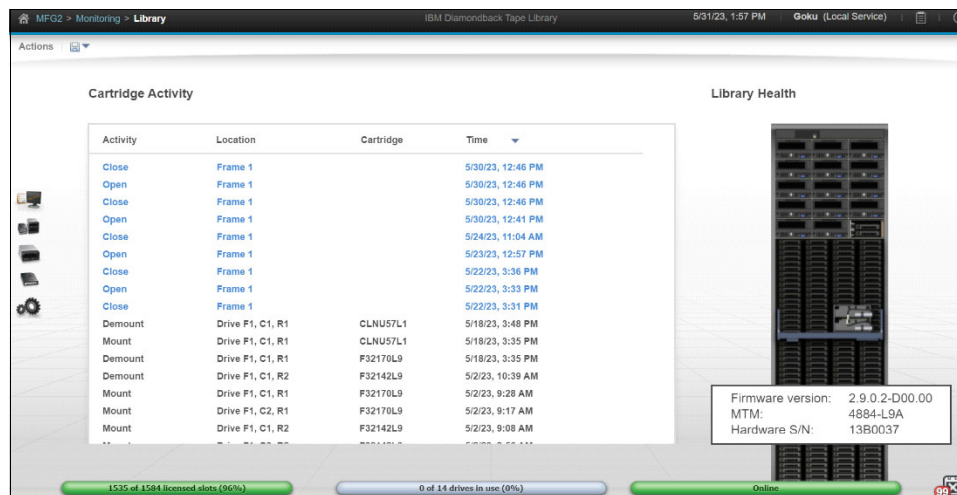


Figure 5-7 Library system page

Viewing the activity log

The recent Library cartridge activity that includes Type, Location, Cartridge Volser, and Time is shown in the activity log. See Figure 5-8.

Cartridge Activity			
Activity	Location	Cartridge	Time ▼
Export	I/O Slot F1, Upper, R1	F32191L9	5/30/23, 4:25 PM
Export	Slot F1, C1, R4, T0	F32191L9	5/30/23, 4:25 PM
Mount	Drive F1, C1, R4	F32191L9	5/30/23, 1:59 PM
Mount	Drive F1, C1, R3	F32182L9	5/30/23, 1:59 PM
Move	Slot F1, C2, R21, T1	F32161L9	5/29/23, 9:59 PM
Move	Slot F1, C3, R11, T1	F32165L9	5/29/23, 9:59 PM
Move	Slot F1, C3, R11, T1	F32165L9	5/29/23, 9:59 PM
Mount	Drive F1, C3, R3	F32167L9	5/29/23, 9:59 PM
Move	Slot F1, C1, R35, T1	F32167L9	5/29/23, 9:58 PM
Demount	Drive F1, C3, R3	F32175L9	5/29/23, 9:58 PM
Move	Slot F1, C1, R35, T1	F32167L9	5/29/23, 9:58 PM
Mount	Drive F1, C2, R4	F32173L9	5/29/23, 9:56 PM
Move	Slot F1, C1, R35, T1	F32167L9	5/29/23, 9:56 PM
Demount	Drive F1, C2, R4	F32161L9	5/29/23, 9:56 PM
Move	Slot F1, C1, R35, T1	F32167L9	5/29/23, 9:56 PM
Move	Slot F1, C3, R15, T1	F32173L9	5/29/23, 9:56 PM

Figure 5-8 Library cartridge activity log

An activity can include the following types:

Mount	A cartridge was inserted into a drive.
Unmount	A cartridge was removed from a drive.
Import	A cartridge was inserted into the library through an I/O slot.
Export	A cartridge was removed from the library through an I/O slot.
Move	A cartridge was moved from one slot to another.
Open	A library door was opened.
Close	A library door was closed.

Viewing component status

View the status of the library frame, tape drives, I/O stations, and the robotic accessor (accessor) on the System page.

You can see the location of each tape drive and if a cartridge is mounted. See Figure 5-9 on page 60.



Figure 5-9 Library component status

If a problem occurs with any of the components, the component turns RED. Go to the Events page to view the detailed information about the error or warning (5.2.3, “Events” on page 69).

Robotic accessor (Accessor)

A red color light on the robotic 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



The I/O stations are located in the lower right corner of the frame’s tape drive slots section similar to where the magazine is physically located in the frame. If a cartridge is in the magazine, it is shown in the GUI. An empty I/O magazine is shown when there are no cartridges in the magazine.

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 5-2 lists each icon.

Table 5-2 Library status icons

Icon	Description
	The Tasks icon shows how many tasks are running.
	The Alerts icon shows how many alerting events are reported for the library.

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 5-10.

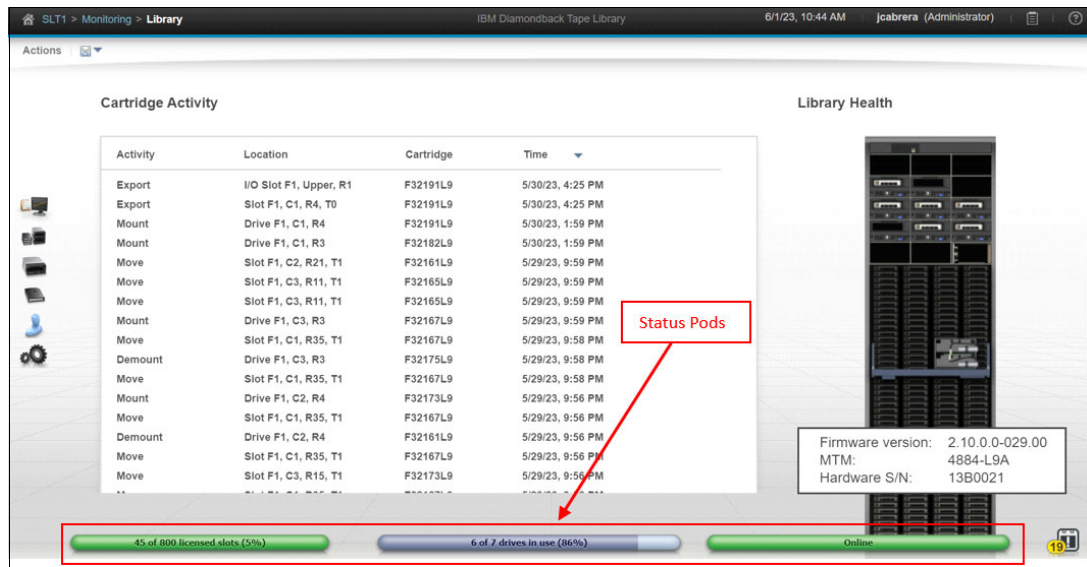


Figure 5-10 Library status pods

Physical capacity status pod

The physical capacity pod displays how many licensed slots are filled by cartridges in this tape library.

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 is shown 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. Hover over the status pod if the pod is yellow or red to see the description 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.

5.2.2 Actions and remote management of the library

Library functions and Actions can be accessed through the menu on the upper-left corner of the library page (see Figure 5-11).

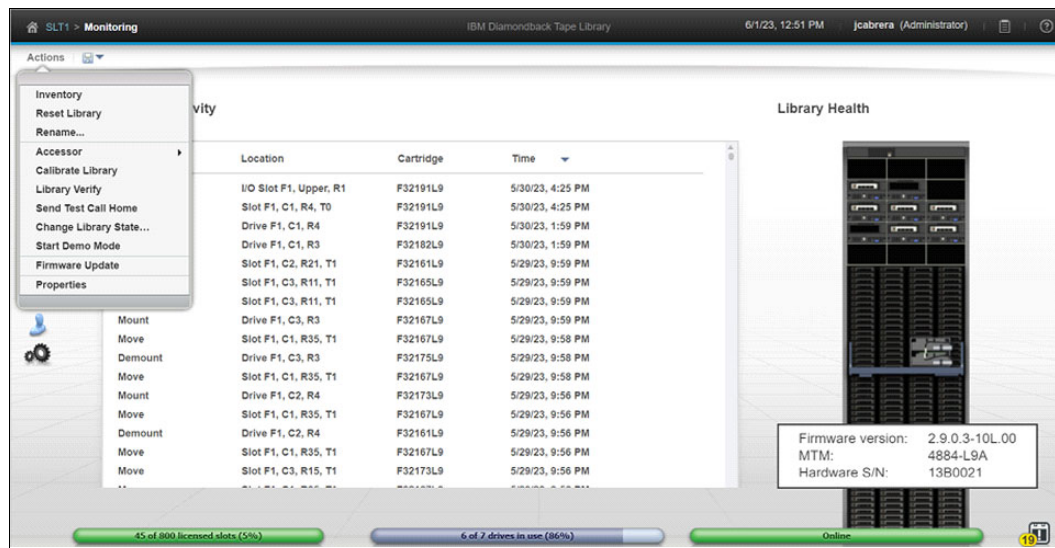


Figure 5-11 library actions menu

Running an inventory

An inventory operation (see Figure 5-12 on page 63) 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.




Figure 5-12 Library inventory command

The IBM Diamondback Tape Library automatically inventories the library. However, certain conditions require that a manual inventory be run. For example, if an automatic inventory of the media was inconsistent, an inventory of the library needs to run to reconcile any inconsistency.

Similarly, a manual inventory may be needed if the host application software indicates that a problem exists with the library inventory.

The IBM Diamondback Tape Library provides two inventory options, as shown in Figure 5-13.




Figure 5-13 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 run for many hours depending on the number of cartridges in the library.

Note: The depth of a cartridge location in an HD slot is known as a *tier*. The cartridges that are immediately accessible by the accessor in an HD slot are Tier 1 cartridges. Behind that is Tier 2, Tier 3, and so on. The maximum tier in an LTO HD slot is Tier 9.

You can inventory the frame by selecting it and then clicking **Actions Inventory**.

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 labels 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 IBM Diamondback management GUI.

Note: It is a best practice to 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 IBM Diamondback 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
- ▶ When this option is used, it is performed concurrently with other move commands. The option can take a maximum of 4 hours to perform a full inventory with an audit on a full HD frame, depending on the numbers of cartridges in the library and host usage.

Reset Library

If the IBM Diamondback Tape Library experiences a stopped state, the reset library command resets the control system of the library. This action (see Figure 5-14 on page 65) resets the node cards in the library.

A reset can take up to 3 minutes. A reset might require you to reconnect to the web server after the reset completes.




Figure 5-14 Reset Library action

Rename the library

Rename the library by using the Actions menu and by selecting **Rename**, as shown in Figure 5-15.

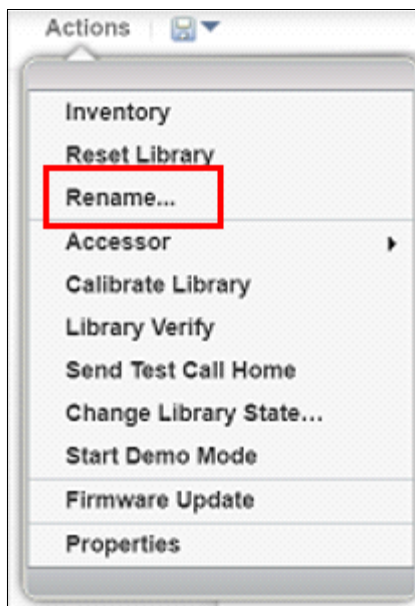


Figure 5-15 Rename library action

Calibrate Library

Use the **Calibrate Library** option to ensure puts, gets, mounts, and dismounts are done accurately into slots, I/O stations, and drives. This task ensures that a proper calibration is performed for slot placement, and a calibration on all the I/O stations and drives in the IBM Diamondback Tape Library. The calibrate library action is for use by service or in a response to instructions included in an events 'fix procedures'.

Library Verify

The **Library Verify** option performs a test motion of the robotic accessor to ensure this is working properly. There is also the option to perform a drive test by mounting a diagnostic cartridge and running read/write test.

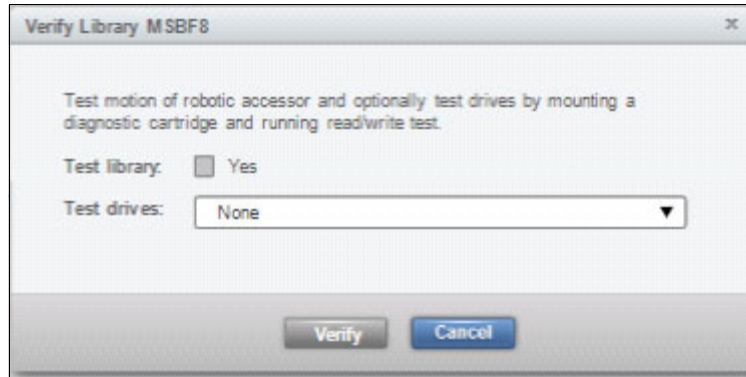


Figure 5-16 Verify library options

Send Test Call Home

Select **Send Test Call Home** to send a notification to IBM by using the TSSC. Contact IBM support to verify that the notification was received.

Change Library State

If you select **Change Library State**, the library can be put into a service state to appear as “In service mode (offline)” to the host. When the library is in service mode, it is not able to accept any request through control path drives including cartridge movement requests.

Start Demo

Selecting **Start Demo Mode** in the library causes the library to perform random moves of cartridges within the IBM Diamondback library. Use this option to demonstrate that the cartridge movements and robotic accessor are working properly. After this activity is done, the cartridges might have a different position within the library inventory.

Updating library firmware

Use the **Firmware Update** option on the Actions menu to nondisruptively update the library firmware and apply the latest library code enhancements.

This command can be used to update the IBM Diamondback Tape Library library firmware without scheduling downtime and without interrupting the job flow or reducing productivity.

The Diamondback 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, such as 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 [Fix Central](#) and install the package on your workstation.

1. Complete the following steps to update the firmware:

From the System page of the IBM Diamondback management GUI, select **Actions Firmware Update**, as shown in Figure 5-17.

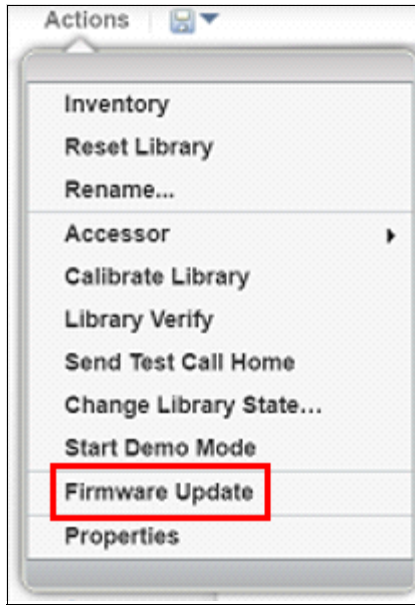


Figure 5-17 Firmware update action

2. On the File Upload window, browse to the library firmware image and click **Open**, as shown in Figure 5-18.




Figure 5-18 Choosing a firmware file for update

Note: The library firmware image file is named `Diamondback_WXYZ.afwz`, where *W*, *X*, *Y*, and *Z* are different numbers or letters, depending on the firmware release version.

3. After it is selected, the Apply Library Firmware window opens. Confirm that the correct file was selected and click **Yes** to continue, as shown in Figure 5-19 on page 68.

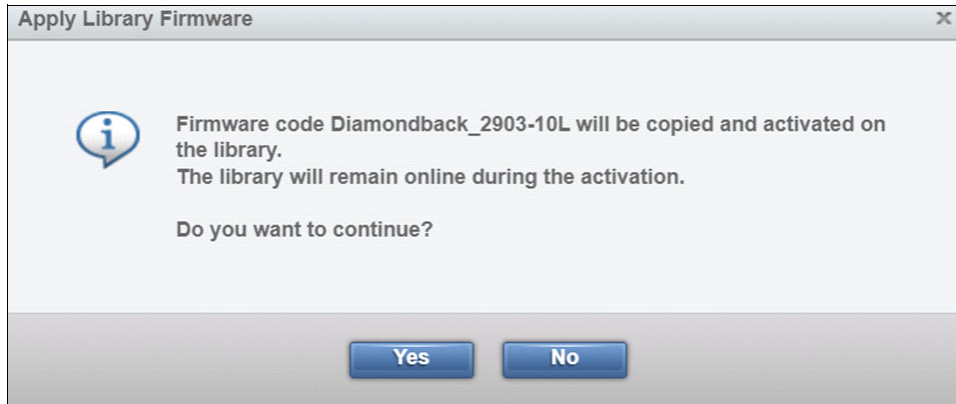


Figure 5-19 Firmware update confirmation

- The firmware downloads to the IBM Diamondback Tape Library. This download can take a few minutes. When finished, 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 5-20.




Figure 5-20 Firmware update status

- The connection to the library will be lost after the firmware is upgraded. Refresh and log back in to the IBM Diamond back Tape Library GUI, then select **Actions Events** to display the status of the update.

Note: If the library firmware update stops, issue a nondisruptive reset of the library from the System page by selecting **Actions Reset Library**. Then, try the firmware update again.

Downloading the library 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 5-21.

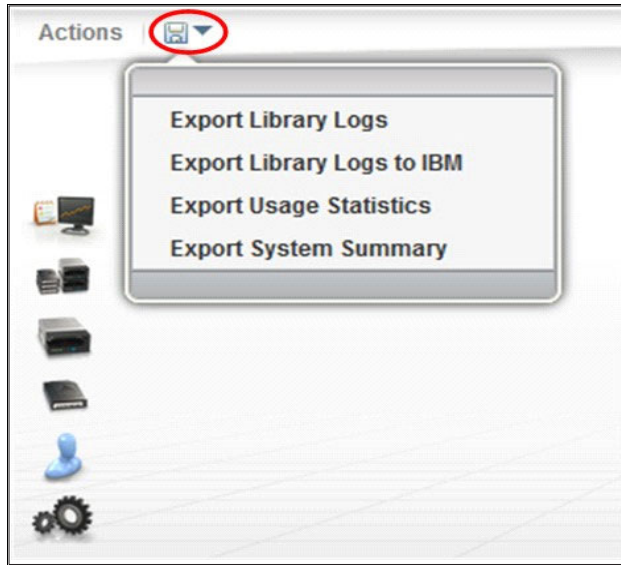


Figure 5-21 Export library logs to files

The following options are available for log 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 by using the Call home function on the 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.

5.2.3 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 5-22.



Figure 5-22 Error icon

Error events are the highest priority events. They require immediate intervention. Errors 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 5-23.



Figure 5-23 Warning icon

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 the library continues to operate smoothly.

Information

The information icon is a blue balloon that contains the lowercase letter i, as shown in Figure 5-24.



Figure 5-24 Information icon

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 user logins 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 event is deleted first. The events table can be limited to show events for 1 day, 1 week, or all. Figure 5-25 on page 71 shows the Events table and filtering option.

Type	Time	User	Description	State
Audit	6/8/23, 4:29 PM	jcabrera	A user logged in to the GUI from 9.86.126.112	
Audit	6/8/23, 11:54 AM	admin	A user logged out of the GUI from 9.67.74.133	
Audit	6/8/23, 10:52 AM	admin	A user logged in to the GUI from 9.67.74.133	
Audit	6/7/23, 5:31 PM	admin	A user logged out of the GUI from 9.86.127.168	
Drive	6/7/23, 4:53 PM	System	The state of drive F1, C1, R2 changed from initializing to onli...	
Audit	6/7/23, 4:53 PM	Service	A drive serial number changed	
Library	6/7/23, 4:53 PM	System	The state of the library changed from driveDegraded to online.	
Drive	6/7/23, 4:53 PM	System	The state of drive F1, C1, R2 changed from unreachable to in...	
Library	6/7/23, 4:53 PM	System	The state of the library changed from online to driveDegraded .	
Drive	6/7/23, 4:53 PM	System	The state of drive F1, C1, R2 changed from initializing to unr...	

Figure 5-25 Library events in the library management GUI

Active events, which are displayed in color within 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 5-26.

The screenshot shows the 'Actions' menu with options: Display Fix Procedure, Mark Active, Mark Inactive, and Properties. The 'Event Properties' dialog is open, showing details for a 'Library' event: Event Type: Library, Location: Accessor 0, Time: Wednesday, April 30, 2014, Mountain Standard Time, User: Service, Description: An obstacle was detected in the accessor or gripper's path, State: closed. The 'Library Fix Procedure' dialog is also open, showing a step: 1) Contact IBM Support. The background shows a list of events with the 'Library' event selected and highlighted in blue.

Figure 5-26 Actions menu event details and fix suggestions

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.




Figure 5-28 Task monitoring action and task details

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 5-29.

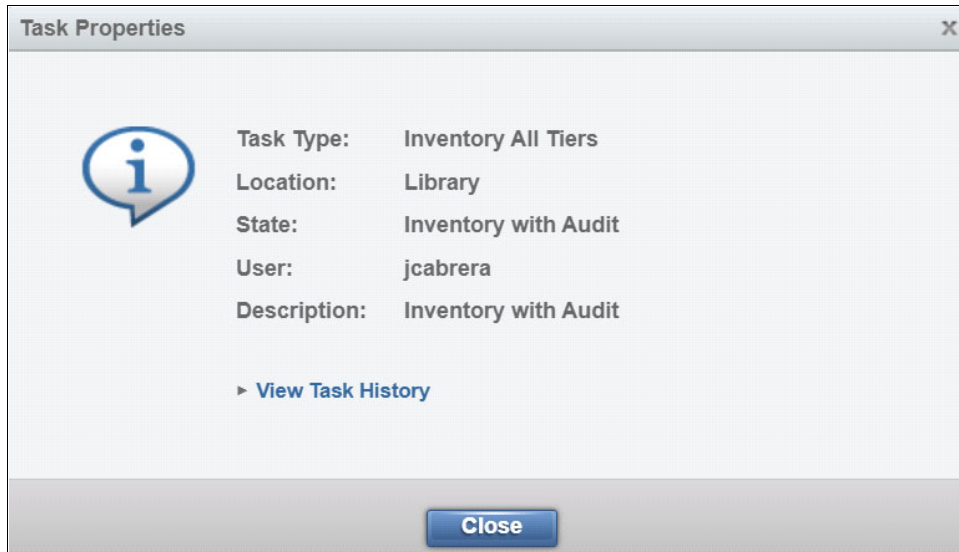


Figure 5-29 Task properties window

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.

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 running tasks. If no icon is displayed, no tasks are active. Hover over the Tasks pod to see the three oldest running tasks sorted in the order that they were created. See Figure 4-174.

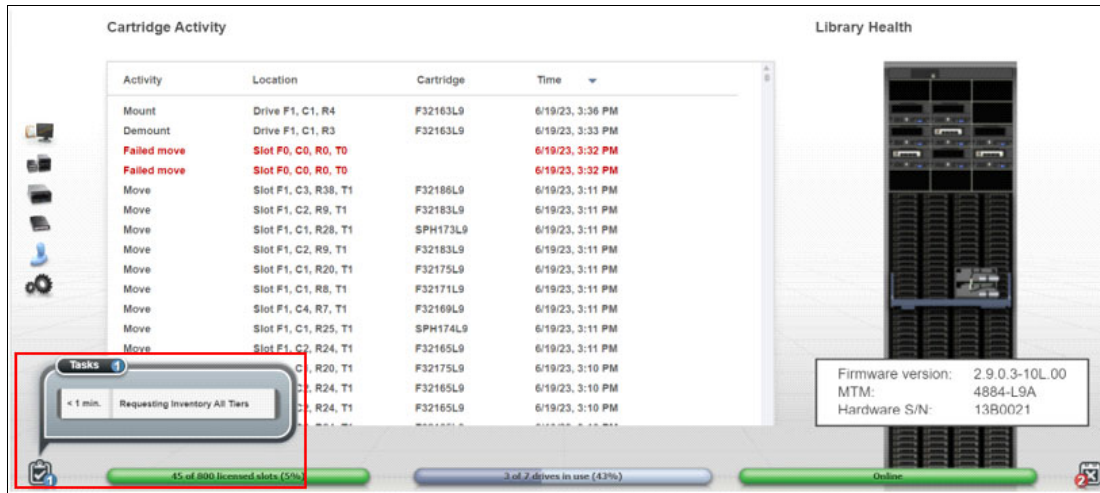


Figure 5-30 Tasks pod

5.3 Library menu

Use the Diamondback management GUI's Library icon to view or manage the logical libraries, as shown in Figure 5-31.



Figure 5-31 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.

Advanced Library Management System

Advanced Library Management System (ALMS) is an extension of the IBM patented Multi-Path Architecture. With ALMS, the IBM Diamondback 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 Diamondback includes ALMS and the ability to immediately start to assign tape drives to any logical library by using the management GUI. Logical libraries can also be added, deleted, or easily changed without disruption. Storage capacity can be changed without affecting the host applications.

ALMS offers the dynamic management of cartridges, cartridge storage slots, tape drives, and logical libraries. It enables the Diamondback 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 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 IBM Diamondback 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 IBM Diamondback 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 are 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 IBM Diamondback 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 IBM Diamondback Tape Library, 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 5-32. 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.

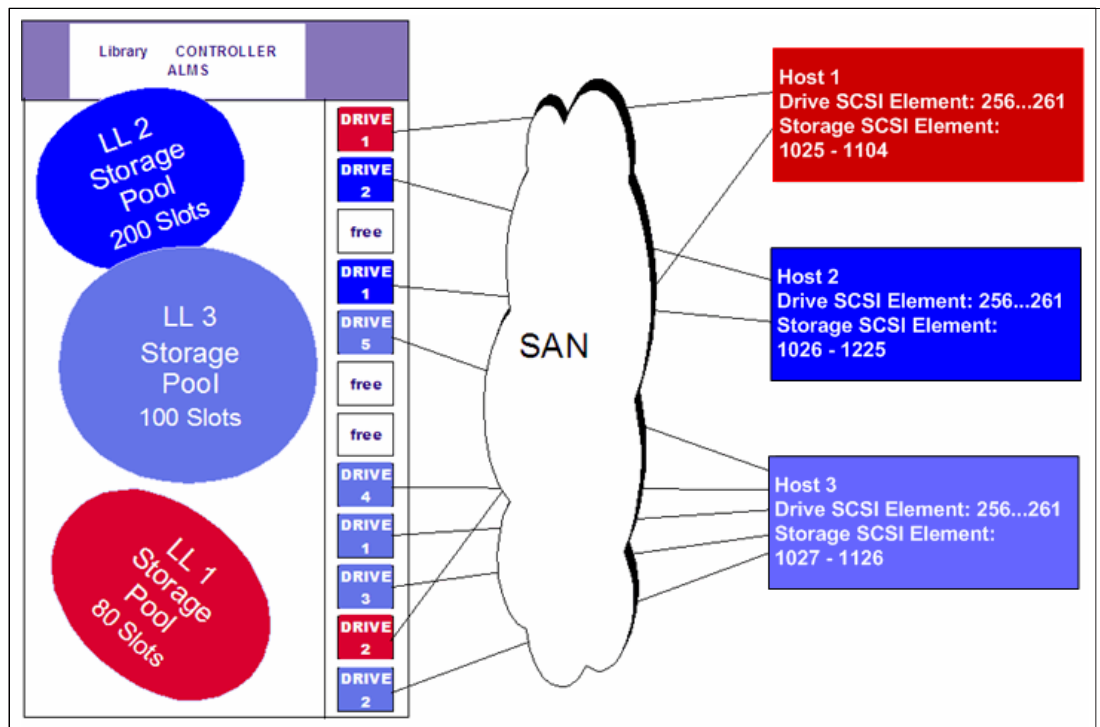


Figure 5-32 IBM Diamondback tape library with ALMS

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 IBM Diamondback 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 IBM Diamondback 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 IBM Diamondback 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 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 IBM Diamondback 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 can make the process of adding and removing cartridges easier and faster.

The Diamondback library has I/O stations with a cartridge magazine so that you can import and export up to 10 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 255 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 and 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 IBM Diamondback Tape Library can be configured into one or more logical libraries that can be shared by multiple applications.

The guided setup for the IBM Diamondback 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 IBM Diamondback 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. The Diamondback 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 IBM Diamondback management GUI to perform the following actions:

Separating 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 5-33 on page 79) requires no special capabilities from the server or application.

Any single logical library (including the entire physical library) can be configured 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

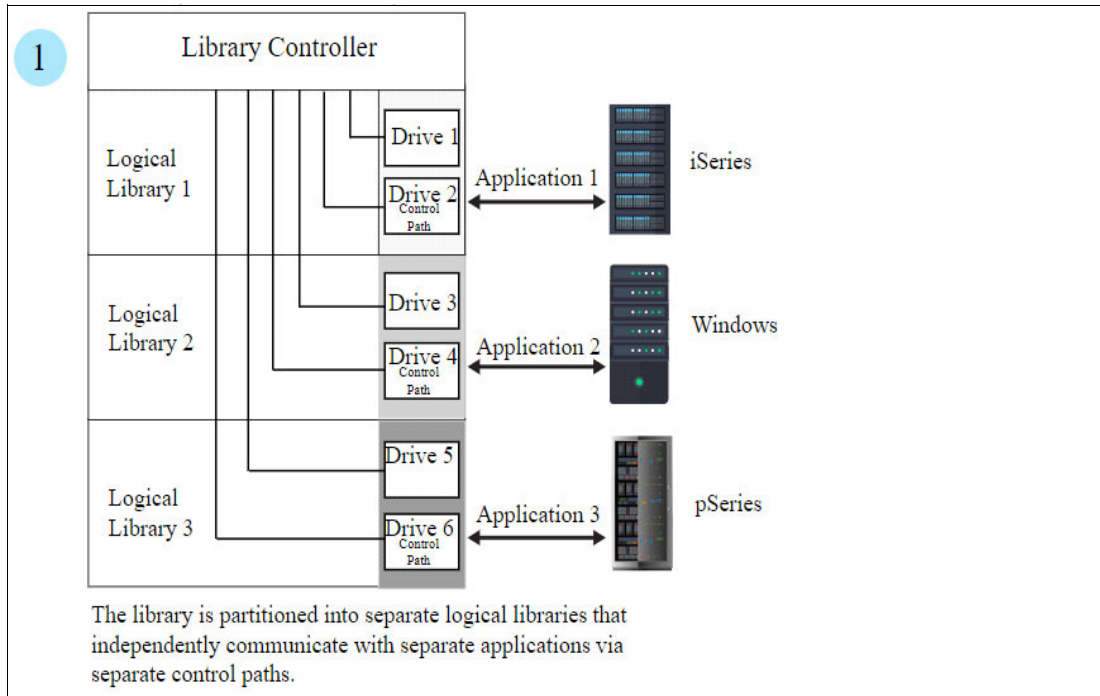


Figure 5-33 Separate logical libraries that are assigned to applications

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 5-34). This configuration is used by IBM Storage Protect.

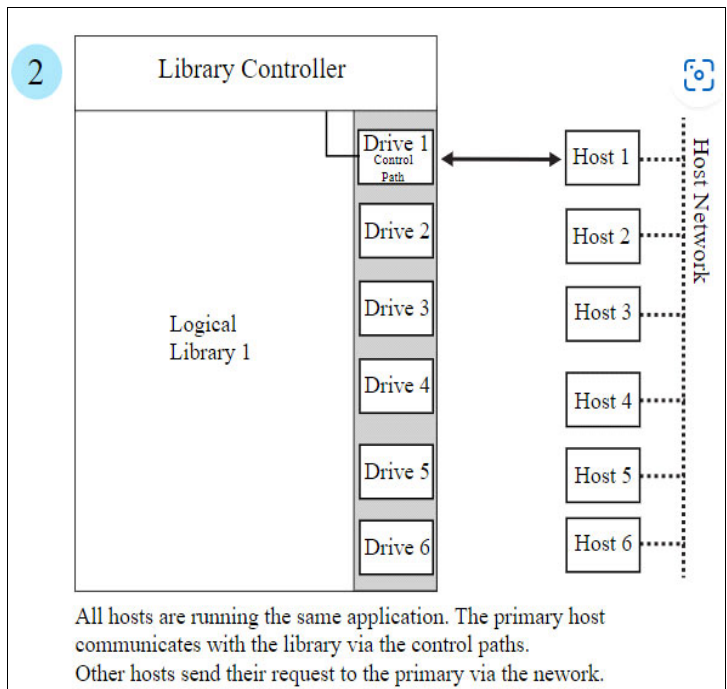


Figure 5-34 Logical libraries to separate applications example

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 5-35). 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.

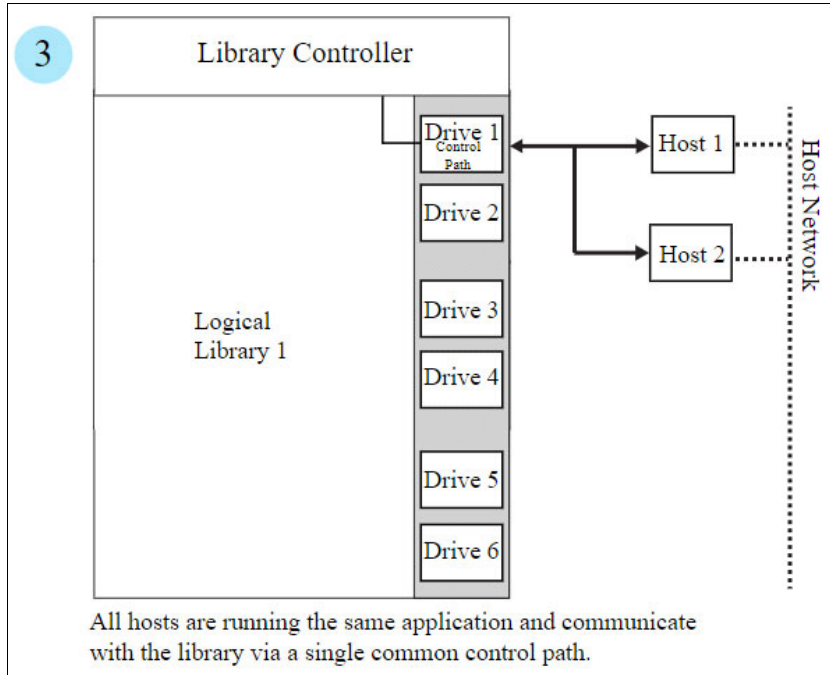


Figure 5-35 Logical libraries to separate applications example

Multiple servers through multiple control paths

Configure multiple logical libraries to communicate with multiple servers through multiple control paths. This configuration (Figure 5-36 on page 81) requires that you add multiple logical libraries, and it is used by separate Backup, Recovery, and Media Services.

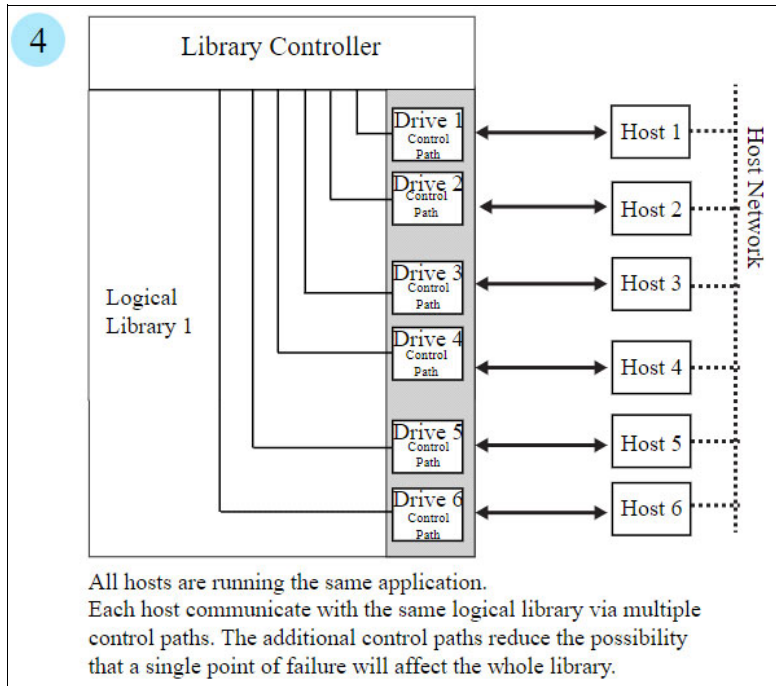


Figure 5-36 Logical libraries to separate applications example

The Diamondback library configuration is not limited to these examples. Many configurations are possible and you can design them according to your business needs.

5.3.1 Creating a logical library

Use the Create Logical Library page to manage and reconfigure logical libraries, including modifying the method of encryption that is used on the logical library. Figure 5-37 shows a list of configured libraries and the window that is used to set the columns of information to display.

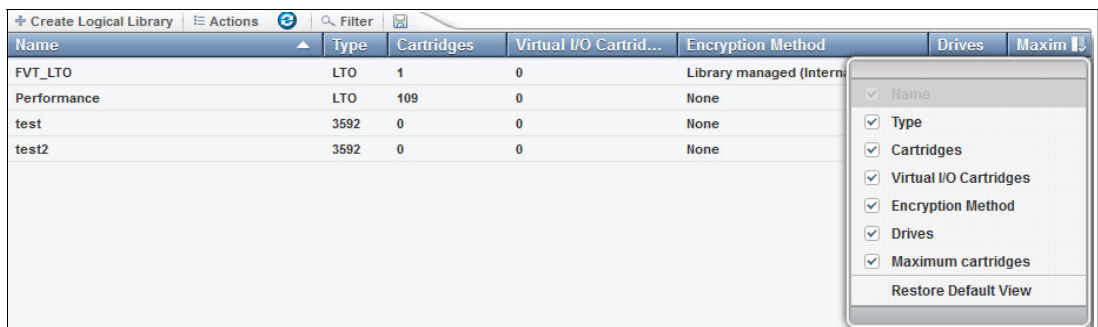


Figure 5-37 Logical library main page

Using multiple logical libraries

You can create multiple libraries within the Diamondback physical library by partitioning the 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. 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 an IBM Diamondback tape library, each logical library shares empty storage slots, empty I/O slots, and the cartridge accessor and uses them 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 5-38.

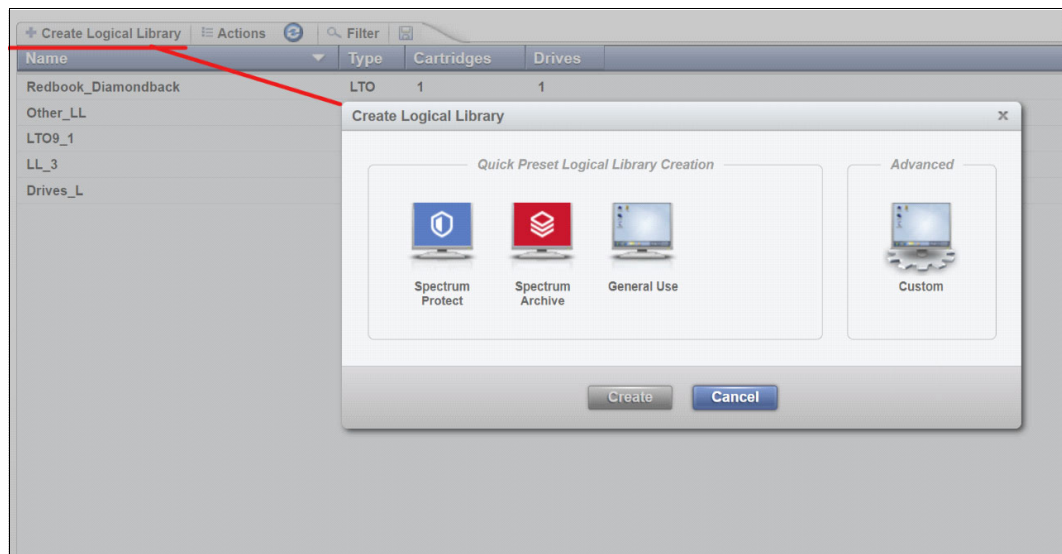


Figure 5-38 Create Logical Library page

IBM Storage Protect

Use this option if you manage your library with IBM Storage Protect (formerly IBM Spectrum Protect).

When you select this option, 2 of the drives that are assigned to the logical library are designated as control path drives. This selection configures all unassigned drives to this logical library.

To configure a logical library by using the IBM Storage Protect option, select the **IBM Spectrum Protect** icon, then define a name, then 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 5-39.

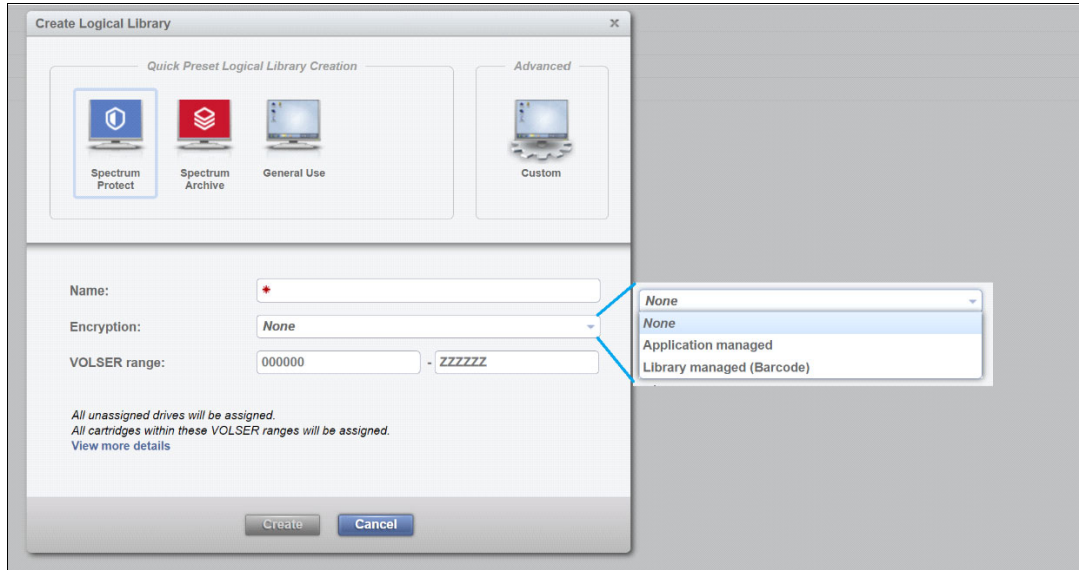


Figure 5-39 Logical library with IBM Storage Protect (IBM Spectrum Protect)

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 147. Also, the key server can be added by selecting **Modify Encryption Method** from the Actions menu, as described in “Modify Encryption Method option” on page 88.

To change drive assignments or control paths, see 5.4.1, “Drives” on page 91.

IBM Storage Archive

Use this option if you plan to use IBM Storage Archive, to incorporate 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 of the drives that are assigned to the logical library are designated as control path drives. 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, 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 5-40 on page 84.

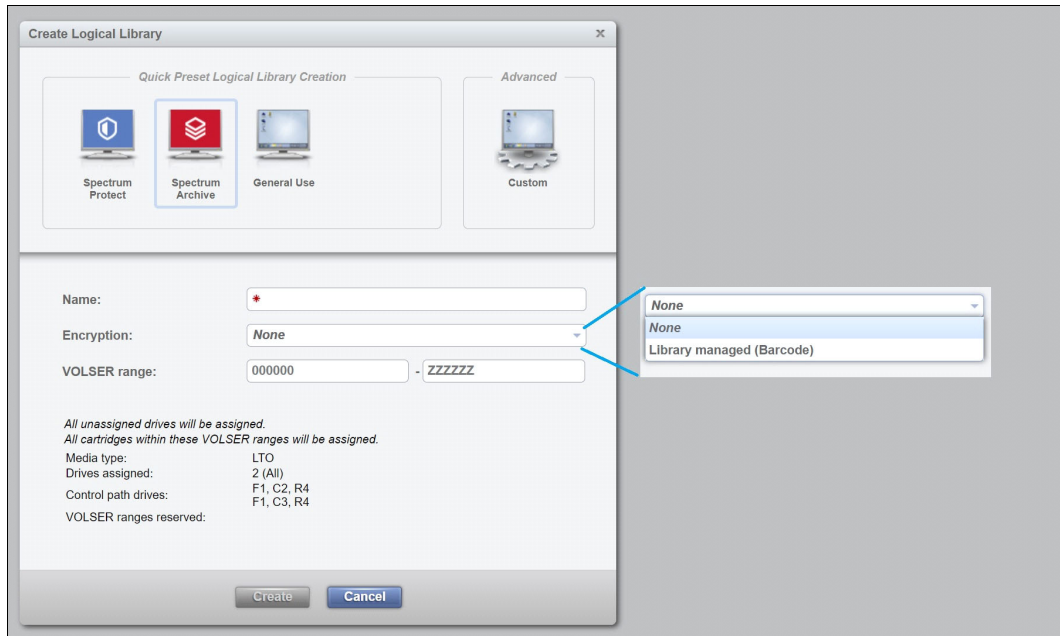


Figure 5-40 Logical library with IBM Storage (Spectrum) Archive

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 147. 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 88.

For more information about changing drive assignments or control paths, see 45.4.1, “Drives” on page 91.

General Use

Select **General Use**, if you have another application that manages your IBM Diamondback 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, 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 5-41 on page 85.

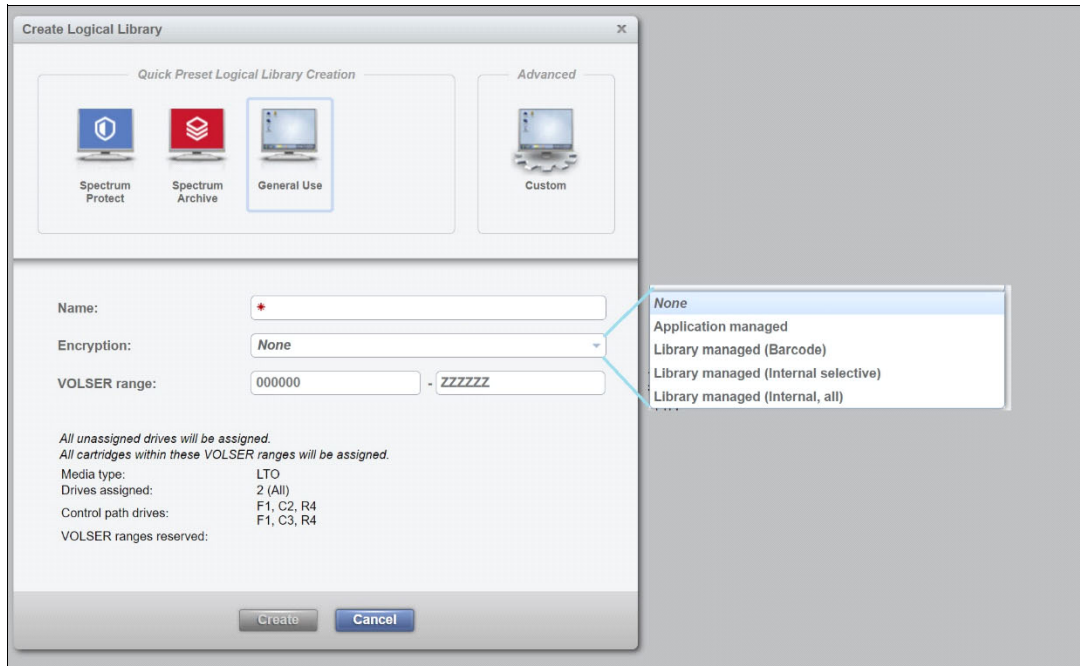


Figure 5-41 Logical library with General Use option settings

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 5-42 on page 86 shows the initial menu after you select **Custom**. Additionally, you can select the drive assignments, VOLSER range, maximum cartridges, and summary.
2. Use the first window to define a name, and select the encryption method to use, if required.

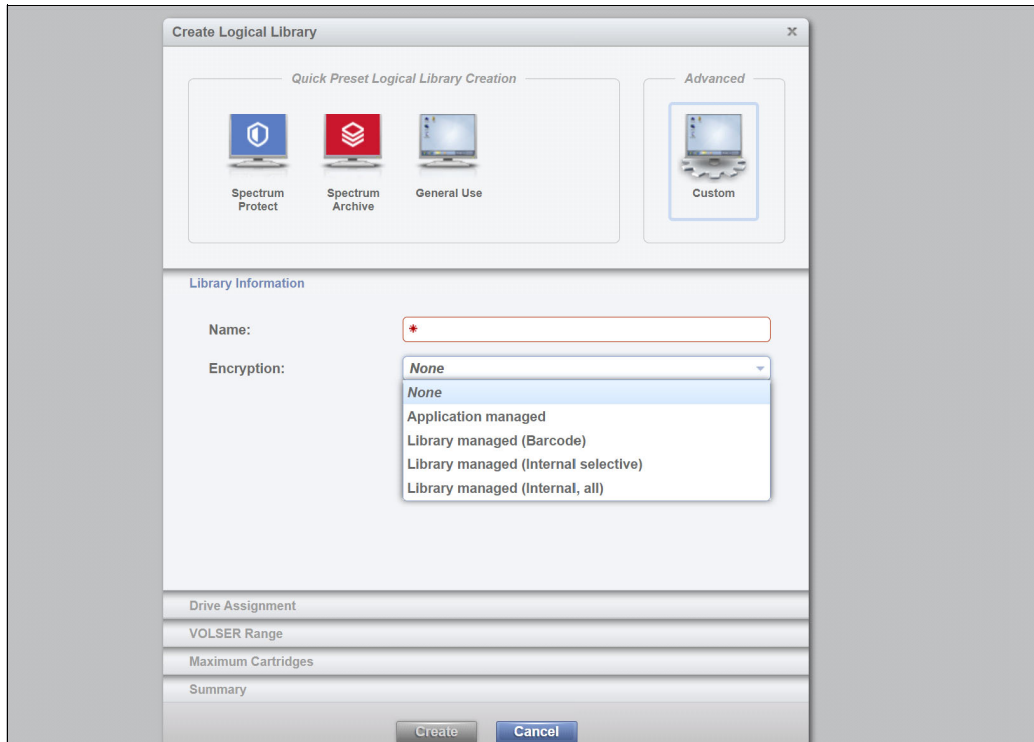


Figure 5-42 Custom Logical library settings

3. Use the Drive Assignment page to configure drives for the logical library, as shown in Figure 5-43. You can select all unassigned drives or select only the drives that you need.

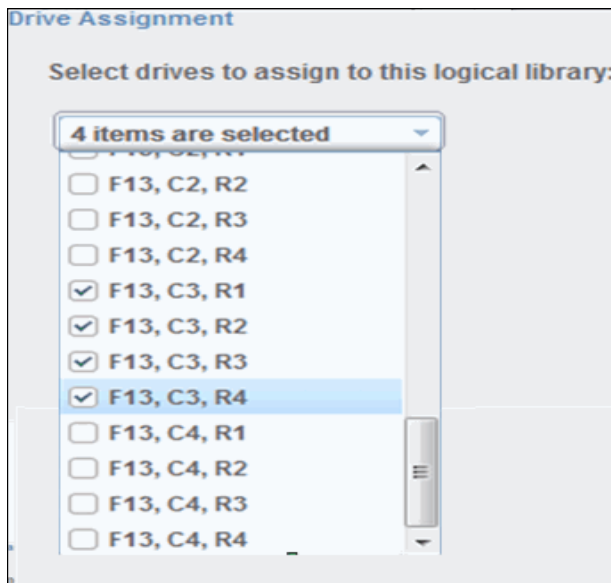


Figure 5-43 Custom drive settings

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 5-44.

VOLSER Range

Identifier -

Apply these range settings only to newly inserted cartridges.

Figure 5-44 Custom VOLSER range settings

5. Use the Maximum Cartridges page to set the number of cartridge slots that are 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 Figure 5-45. The system default is the total number of slots that is licensed in the library.

Maximum Cartridges

Use system default:

Maximum cartridges:

Figure 5-45 Custom Maximum cartridges settings

6. Use the Summary window to display all of your settings, as shown in Figure 5-46.

Summary

Name:

Media type: LTO

Encryption: None

Drives assigned: 2 (All)

Control path drives: F1, C2, R4
F1, C3, R4

Max cartridges: 100

VOLSER ranges reserved: AAA111 - AAA199

Create Cancel

Figure 5-46 Summary example

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 92.

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 92.

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” on page 77. 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 115.

Note: Cleaning cartridges are shared among the logical libraries.

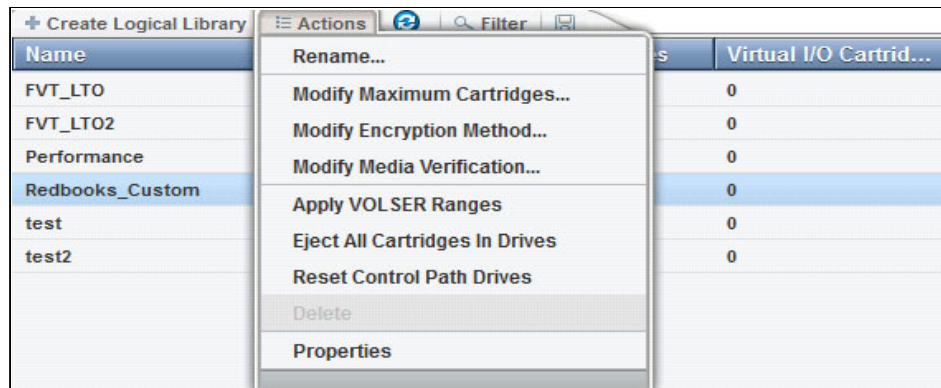


Figure 5-47 Logical library actions menu

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, 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.

Modify Encryption Method option

Use the **Modify Encryption Method** option in 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 to set up or remove encryption from the logical library. With library-managed encryption, 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.

Add up to four servers to the logical library, as shown on Figure 5-48. This page can be used to add a new unlisted EKM server IP address and to test it by clicking **Ping**.

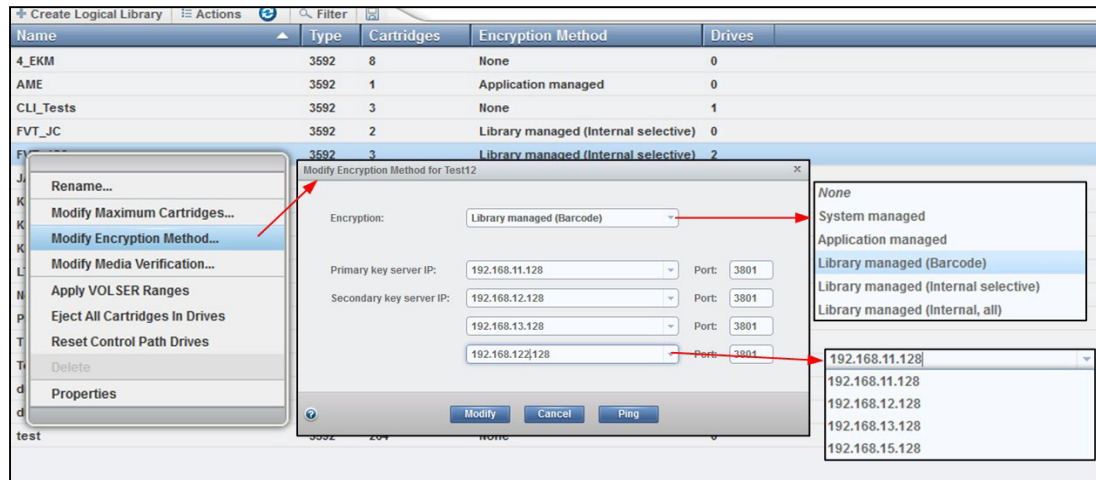


Figure 5-48 Modifying the encryption method

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 IBM Diamondback for as long as 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 93)
- ▶ Delete all VOLSER ranges (see “Deleting a VOLSER range” on page 116)
- ▶ Unassign any cartridges that are assigned (see “Unassigning cartridges” on page 107)

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 5-49.

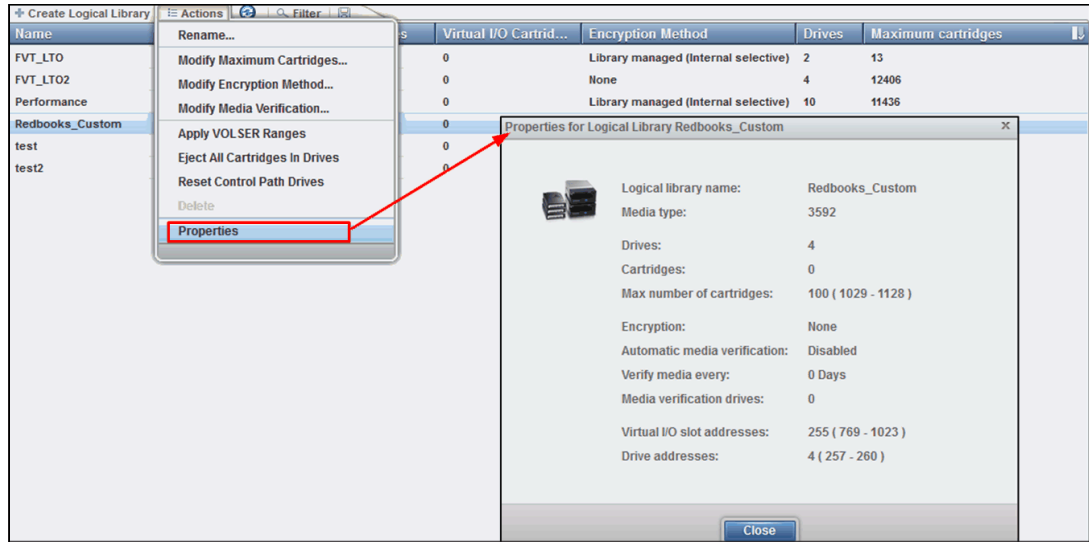


Figure 5-49 Logical library properties

5.3.2 Slots

The Slots option is available for any 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 column, row, and slot number from the menu, as shown in Figure 5-50, and then select either the Set Online action or the **Set Offline** action.

Single or multiple slots can be selected with the number of slots selected being displayed in the upper-right corner of the menu.



Figure 5-50 Library slots page

5.4 Drives menu

In the Diamondback management GUI pages, hover over the **Drives** icon (see Figure 5-51) to view all drives in the library, view drives by logical library, view drive port status. The pages can also be used to perform tasks, such as assigning drives, enabling control paths, configuring drives for verification, and setting drive port configuration.



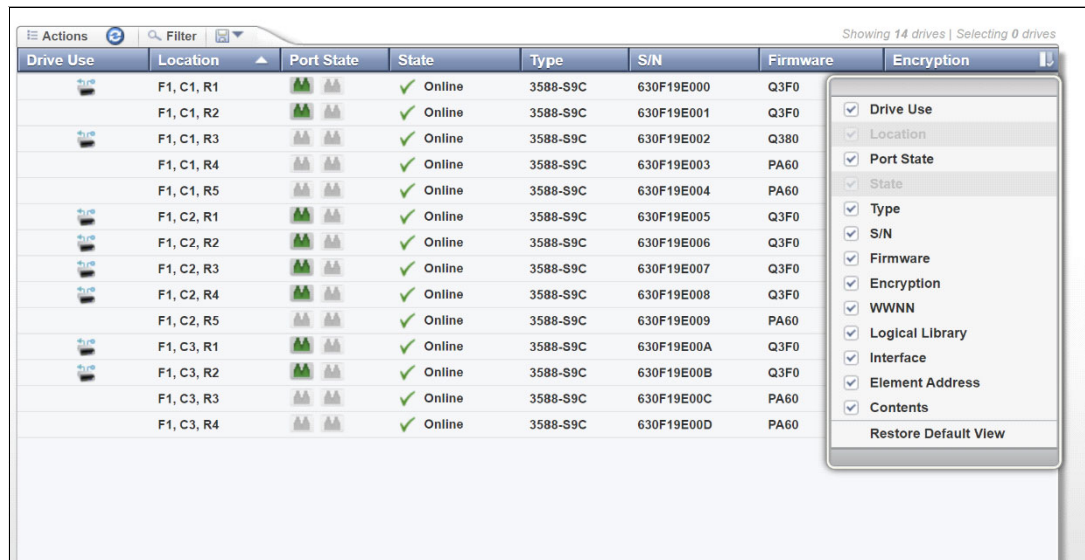
Figure 5-51 Drives menu

5.4.1 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.

The Drives page also contains options and displays additional drive status information, including port status, drive serial numbers (S/Ns), and worldwide node name (WWNN).

Starting with R9 FP1 code a new “Operation” column has been added as shown in Figure 5-52.



Drive Use	Location	Port State	State	Type	S/N	Firmware	Encryption
	F1, C1, R1		✓ Online	3588-S9C	630F19E000	Q3F0	
	F1, C1, R2		✓ Online	3588-S9C	630F19E001	Q3F0	
	F1, C1, R3		✓ Online	3588-S9C	630F19E002	Q380	
	F1, C1, R4		✓ Online	3588-S9C	630F19E003	PA60	
	F1, C1, R5		✓ Online	3588-S9C	630F19E004	PA60	
	F1, C2, R1		✓ Online	3588-S9C	630F19E005	Q3F0	
	F1, C2, R2		✓ Online	3588-S9C	630F19E006	Q3F0	
	F1, C2, R3		✓ Online	3588-S9C	630F19E007	Q3F0	
	F1, C2, R4		✓ Online	3588-S9C	630F19E008	Q3F0	
	F1, C2, R5		✓ Online	3588-S9C	630F19E009	PA60	
	F1, C3, R1		✓ Online	3588-S9C	630F19E00A	Q3F0	
	F1, C3, R2		✓ Online	3588-S9C	630F19E00B	Q3F0	
	F1, C3, R3		✓ Online	3588-S9C	630F19E00C	PA60	
	F1, C3, R4		✓ Online	3588-S9C	630F19E00D	PA60	

Figure 5-52 Drives page

The Drives page and the Actions menu options are shown on Figure 5-53. The drives can be assigned, unassigned, or reassigned to a logical library from this menu. Other drive actions are also performed from this menu.

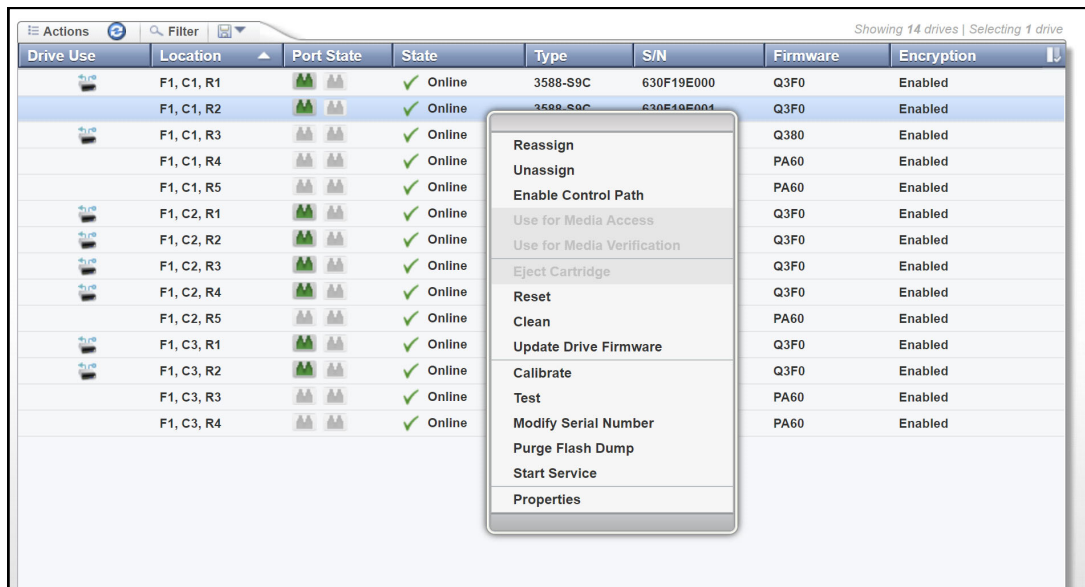


Figure 5-53 Drives page action menu

Adding or removing tape drives

A tape drive is unassigned when it is installed in the tape library. Any unassigned drives are automatically assigned when a logical library is created by using the logical library presets. You can also 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 5.3.1, “Creating a logical library” on page 81.

Reassigning a drive

Reassigning a drive from one logical library to another logical library is easily done by using the Reassign drive option. Ensure that the drive is not a control path drive. Use the Drives by Logical Library page to reassign the drive, as shown in Figure 5-54.

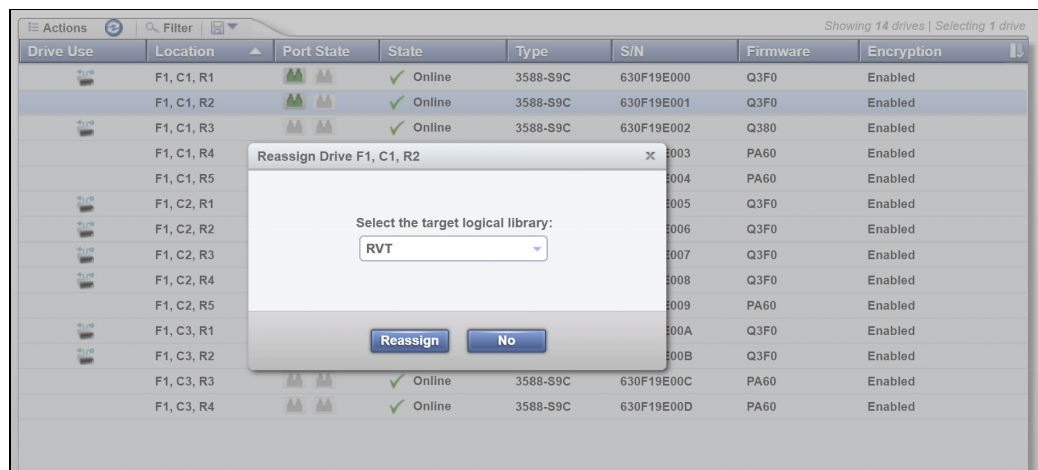


Figure 5-54 Reassign a drive example

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 5-55.

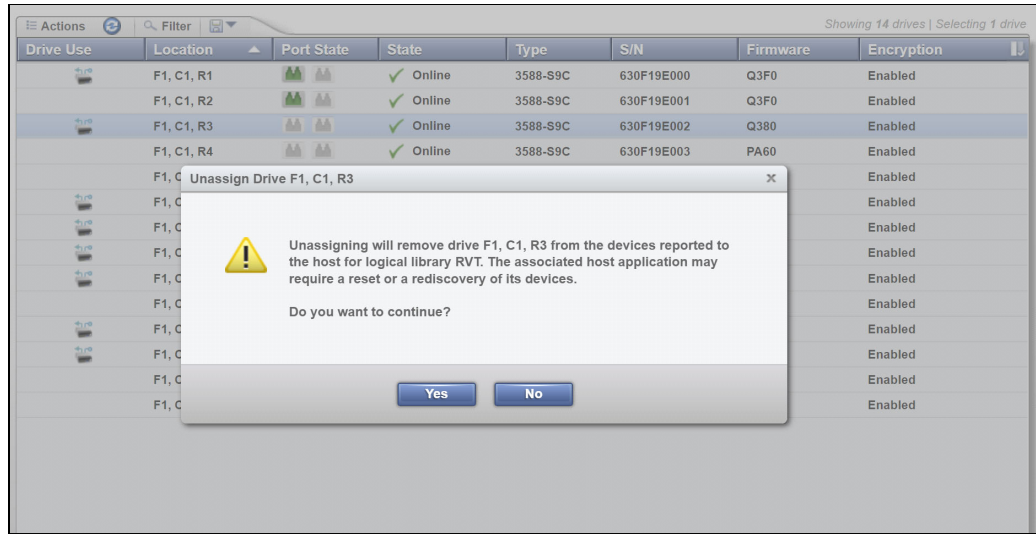


Figure 5-55 Unassign drive example

Control path drives

A *control path* is a logical path to the Diamondback 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 Diamondback 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. This 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 5-56.



Figure 5-56 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. See Figure 5-57 on page 94.

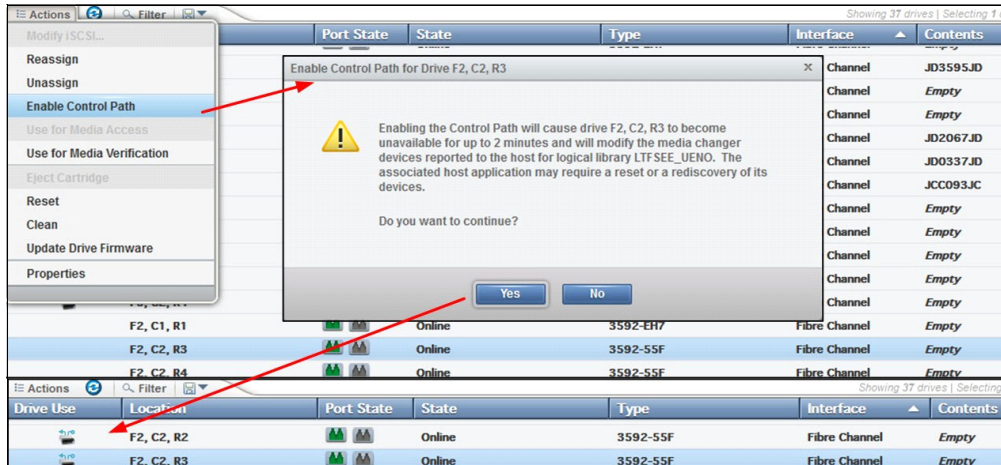


Figure 5-57 Enable control path

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 5-58.

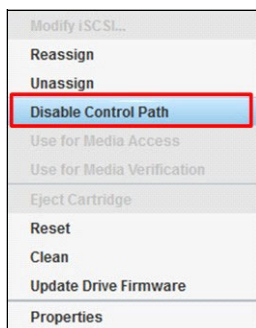


Figure 5-58 Disable control path

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.

Resetting a tape drive

Reset a drive by using the Drives by Logical Library page. Select the drive. Then, select **Actions Reset**, as shown in Figure 5-59. The drive is unavailable for up to 4 minutes.

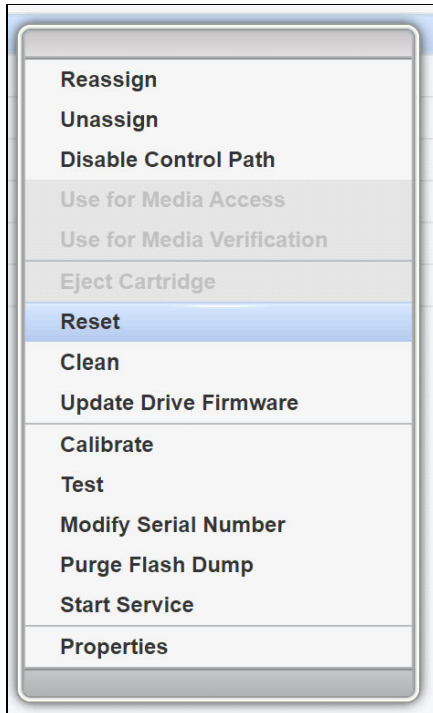


Figure 5-59 Drive reset option

Cleaning tape drives

The Diamondback tape library automatically cleans tape drives as needed. However, you can also start a manual cleaning from the Drives page (see Figure 5-60). Select **Actions Clean**.

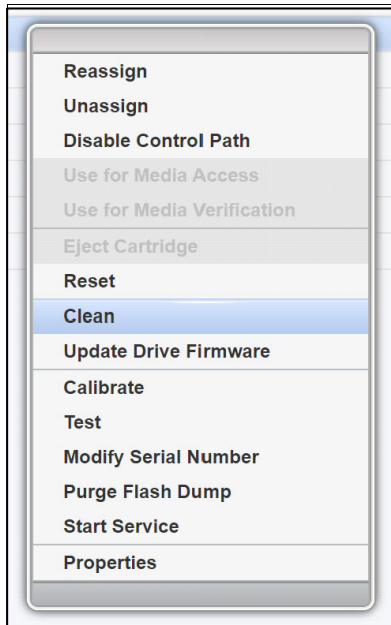


Figure 5-60 Clean drive option

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 112.

Updating drive firmware

As shown in Figure 5-61, select one drive or multiple drives of the same media type to update simultaneously with the firmware level that was uploaded. A drive is updated even if it is running a firmware level that is a later version of drive code than what was uploaded.

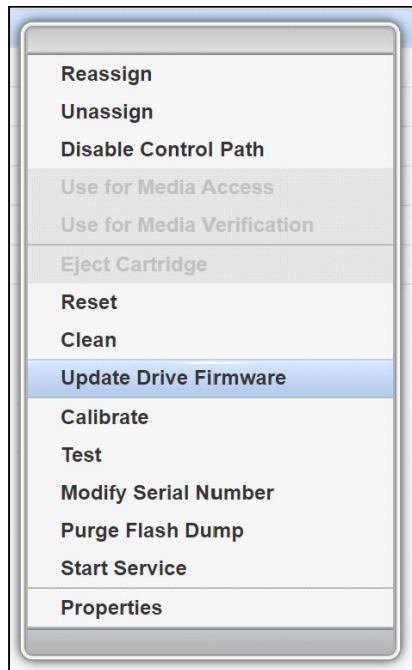


Figure 5-61 Update drive firmware option

Drive firmware can be updated on one drive or multiple drives at the same 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 releases than the version that the drive is currently 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](#).

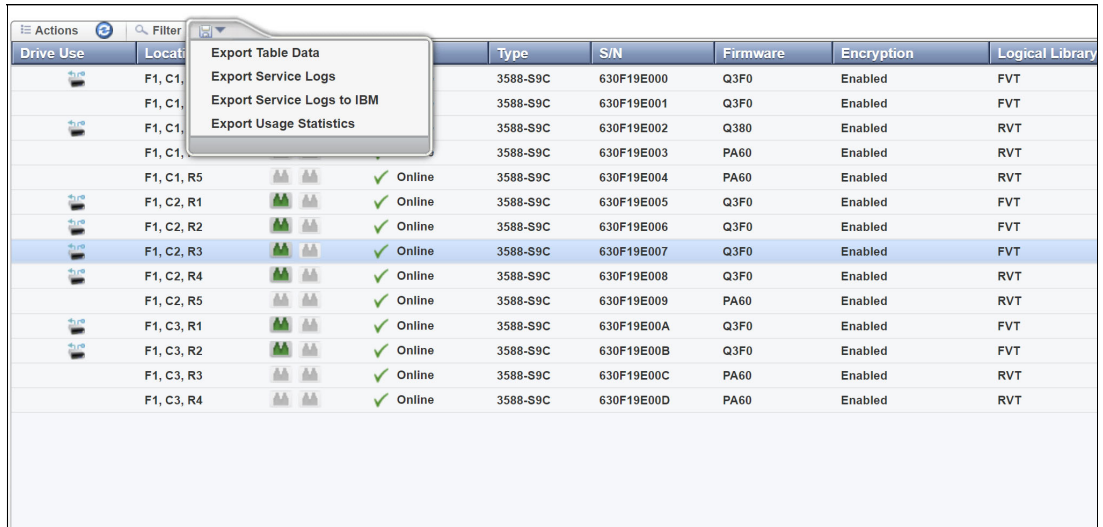
From the Drives page of the Diamondback management GUI, select one or more drives, and click **Actions Update Drive Firmware**, as shown in Figure 5-61. When the File Upload window opens, browse to the drive firmware image and click **Open**.

Note: The LTO9 drive firmware image file has a .fmrz suffix.

Monitor the status of the update in the State column and the Firmware column of the Drives table. After the update finishes, 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 5-62.



Drive Use	Location	Type	S/N	Firmware	Encryption	Logical Library
	F1, C1,	3588-S9C	630F19E000	Q3F0	Enabled	FVT
	F1, C1,	3588-S9C	630F19E001	Q3F0	Enabled	FVT
	F1, C1,	3588-S9C	630F19E002	Q380	Enabled	RVT
	F1, C1,	3588-S9C	630F19E003	PA60	Enabled	RVT
	F1, C1, R5	3588-S9C	630F19E004	PA60	Enabled	RVT
	F1, C2, R1	3588-S9C	630F19E005	Q3F0	Enabled	FVT
	F1, C2, R2	3588-S9C	630F19E006	Q3F0	Enabled	FVT
	F1, C2, R3	3588-S9C	630F19E007	Q3F0	Enabled	FVT
	F1, C2, R4	3588-S9C	630F19E008	Q3F0	Enabled	RVT
	F1, C2, R5	3588-S9C	630F19E009	PA60	Enabled	RVT
	F1, C3, R1	3588-S9C	630F19E00A	Q3F0	Enabled	FVT
	F1, C3, R2	3588-S9C	630F19E00B	Q3F0	Enabled	FVT
	F1, C3, R3	3588-S9C	630F19E00C	PA60	Enabled	RVT
	F1, C3, R4	3588-S9C	630F19E00D	PA60	Enabled	RVT

Figure 5-62 Exporting data, logs or Statistics

When exporting 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.

If needed, use this window to export service logs from a drive to send to IBM Support. Select the drive for which you want to export the logs, and then select **Export Service Logs** from the Export Data menu.

5.4.2 Drives by Logical Library page

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.

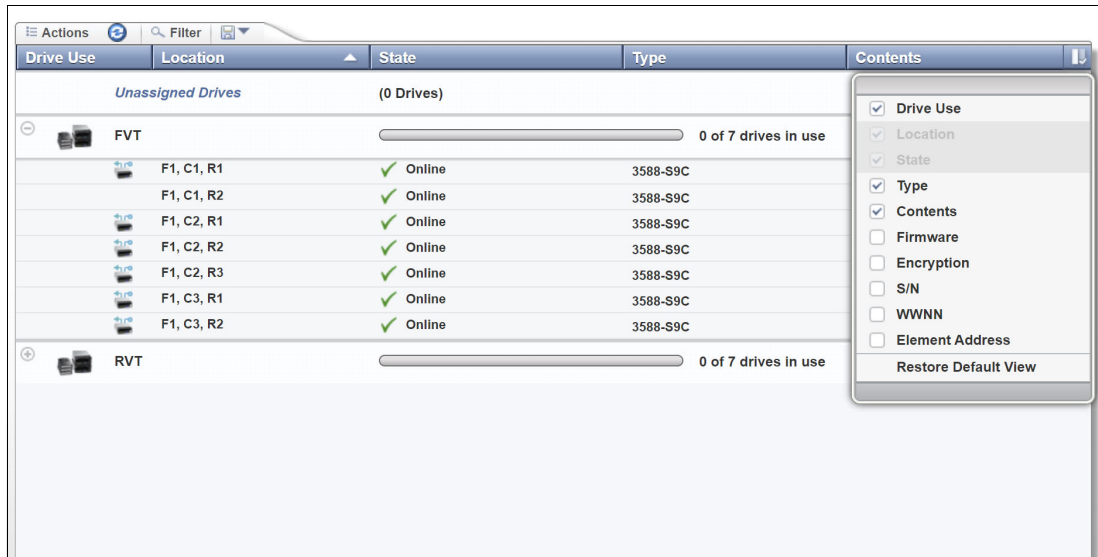


Figure 5-63 Drives by logical library page

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 5-64.

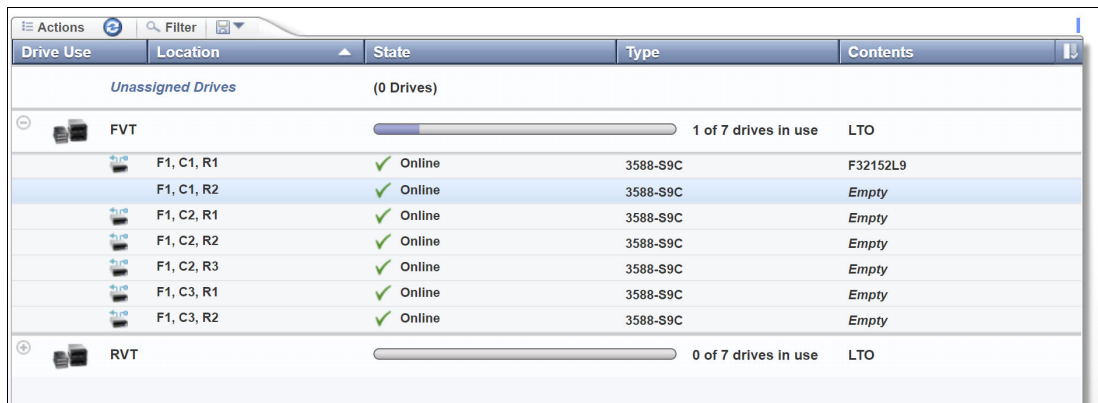


Figure 5-64 Monitoring drive usage

If all of the drives are frequently in use, the logical library might be over-allocated. This can be addressed by assigning additional drives to the logical library.

The same actions can be performed on this page as on the Drives menu. For more information about these functions, see 5.8, “Remote support” on page 151.

5.4.3 Fibre Channel Ports option

Use the Fibre Channel Ports option in the Drives menu (5.8, “Remote support” on page 151) to view all the Fibre Channel port information for 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 5-65 on page 99.









Drive	Port	State	WWPN	Speed	Topology	Loop Id
F1, C1, R1	0		50:05:07:60:78:40:00:F0	Auto (8Gb/s)	Auto-L (N Port)	1
F1, C1, R4	0		50:05:07:60:78:40:00:F3	Auto (8Gb/s)	Auto-L (N Port)	4
F1, C3, R1	0		50:05:07:60:78:40:00:FA	Auto (8Gb/s)	Auto-L (N Port)	9
F1, C2, R4	0		50:05:07:60:78:40:00:F8	Auto (8Gb/s)	Auto-L (N Port)	8
F1, C2, R5	1		50:05:07:60:78:80:00:F9	Auto (8Gb/s)	Auto-L (N Port)	73
F1, C3, R2	1		50:05:07:60:78:80:00:FB	Auto (8Gb/s)	Auto-L (N Port)	74
F1, C1, R1	1		50:05:07:60:78:80:00:F0	Auto ()	Auto-L ()	65
F1, C1, R2	1		50:05:07:60:78:80:00:F1	Auto ()	Auto-L ()	66

Figure 5-65 Drives port summary

The port State column is sorted by state color (green, yellow, or gray). The gray port status icon is shown in Figure 5-66.



Figure 5-66 Gray port status icon

The gray port status icon indicates that no light is detected at the Fibre Channel port of the drive. 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 5-67.



Figure 5-67 Yellow port status icon

The yellow port status icon indicates that the Fibre Channel connections between the drive and the Fibre Channel device, such as a host bus adapter, switch, or hub) to which it is connected are good. The port is logged in to 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 5-68.



Figure 5-68 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 **Actions** → **Modify Topology** from the Ports view, as shown in Figure 5-69.

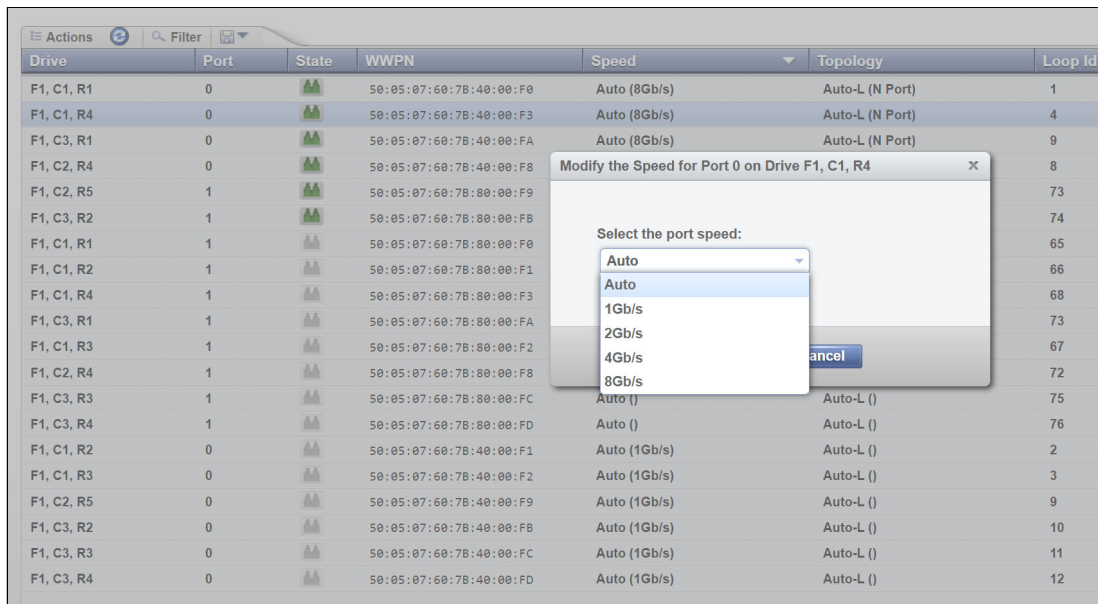


Figure 5-69 Modifying the topology and speed in the Ports view

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

From the Actions menu, you can display the properties of a port, as shown in Figure 5-70 on page 101.

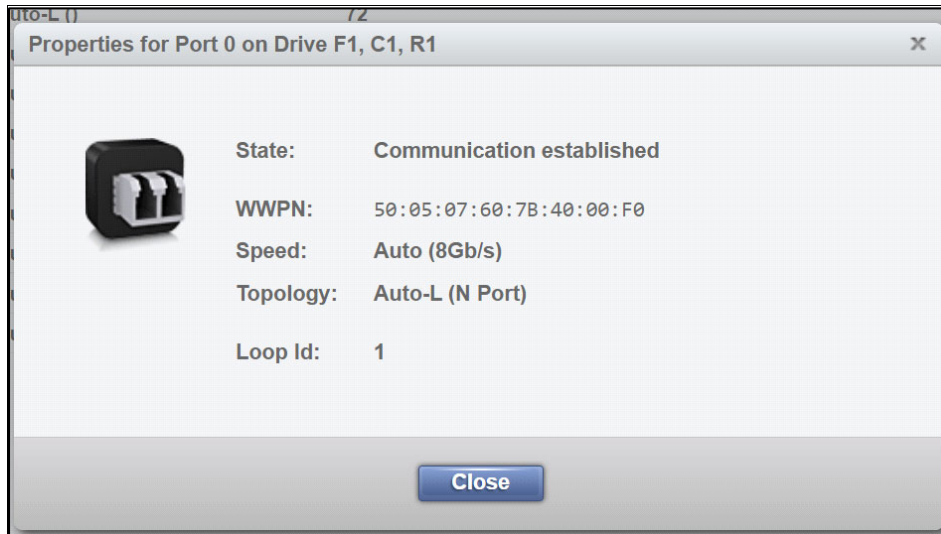


Figure 5-70 Port properties

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 5-71 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.

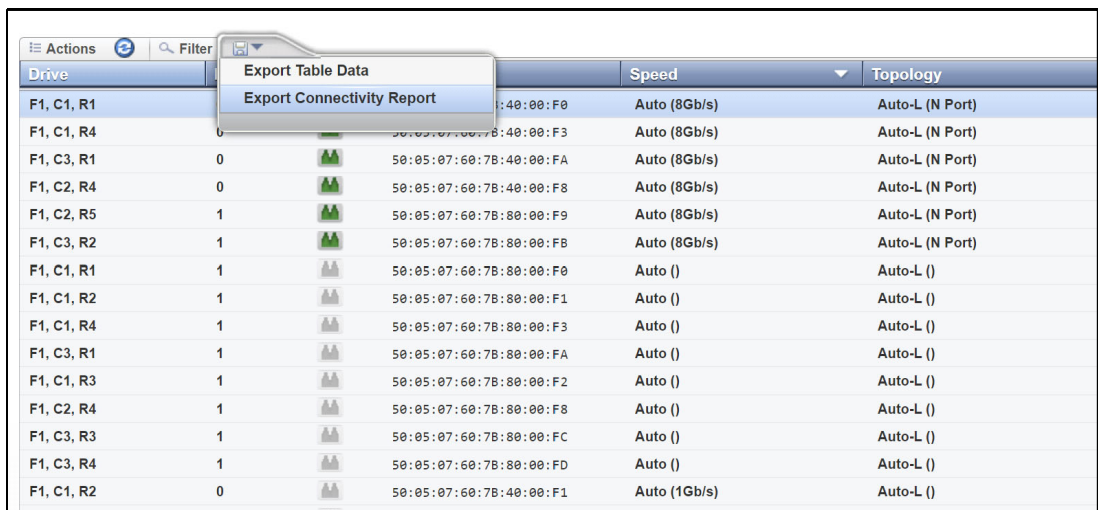


Figure 5-71 Exporting drive connectivity report

5.4.4 SAS Ports

Use the SAS Ports option in the Drives menu (see Figure 5-72 on page 102) to view all 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, C1, R1	0		50:05:07:63:0F:59:E0:00	Auto ()	1
F1, C1, R1	1		50:05:07:63:0F:99:E0:00	Auto (1Gb/s)	65
F1, C1, R2	0		50:05:07:63:0F:59:E0:01	Auto ()	2
F1, C1, R2	1		50:05:07:63:0F:99:E0:01	Auto (1Gb/s)	66
F1, C1, R3	0		50:05:07:63:0F:59:E0:02	Auto (1Gb/s)	3
F1, C1, R3	1		50:05:07:63:0F:99:E0:02	Auto (1Gb/s)	67
F1, C1, R4	0		50:05:07:63:0F:49:E0:03	Auto (1Gb/s)	4
F1, C1, R4	1		50:05:07:63:0F:89:E0:03	Auto ()	68
F1, C1, R5	0		50:05:07:63:0F:59:E0:04	Auto (1Gb/s)	5
F1, C1, R5	1		50:05:07:63:0F:99:E0:04	Auto (1Gb/s)	69
F1, C2, R1	0		50:05:07:63:0F:59:E0:05	Auto ()	5
F1, C2, R1	1		50:05:07:63:0F:99:E0:05	Auto (1Gb/s)	69
F1, C2, R2	0		50:05:07:63:0F:59:E0:06	Auto ()	6
F1, C2, R2	1		50:05:07:63:0F:99:E0:06	Auto (1Gb/s)	70
F1, C2, R3	0		50:05:07:63:0F:59:E0:07	Auto ()	7
F1, C2, R3	1		50:05:07:63:0F:99:E0:07	Auto (1Gb/s)	71

Figure 5-72 SAS ports

For more information about exporting the drive connectivity from the SAS Ports page, see “Exporting drive connectivity report” on page 101.

The port State column is sorted by state color (green, yellow, or gray). The gray port status icon is shown in Figure 5-73.



Figure 5-73 Gray port status icon

The gray port status icon indicates that the SAS port is not initialized.

If no energy is detected at the SAS port of the drive, 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 Figure 5-74.



Figure 5-74 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 5-75 on page 103.



Figure 5-75 SAS port properties

5.5 Cartridges menu

Use the Diamondback 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 5-76.

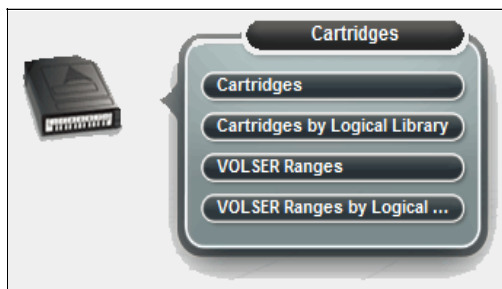


Figure 5-76 Cartridges menu

5.5.1 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)
- ▶ Moving cartridges
- ▶ Exporting the mount history

Figure 5-77 on page 104 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 5-77 Cartridges page

Searching for cartridges

You can use the Cartridges table to find specific cartridges and view their state (Normal, Importing). The table is sortable to help locate specific tape cartridges. You can sort the table by VOLSER, state, location, logical library, element address, or most recent use. The Cartridges by Logical Library page shows the cartridges organized by their logical library, allowing you to easily see which volumes belong to each logical library defined.

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 5-78. If the entire value is unknown, substrings can be used for the search. Wildcard characters are not supported. Instead of wildcard characters, add substrings to the search by clicking the plus sign (+) at the end of each criterion. (Search values are not case-sensitive.)

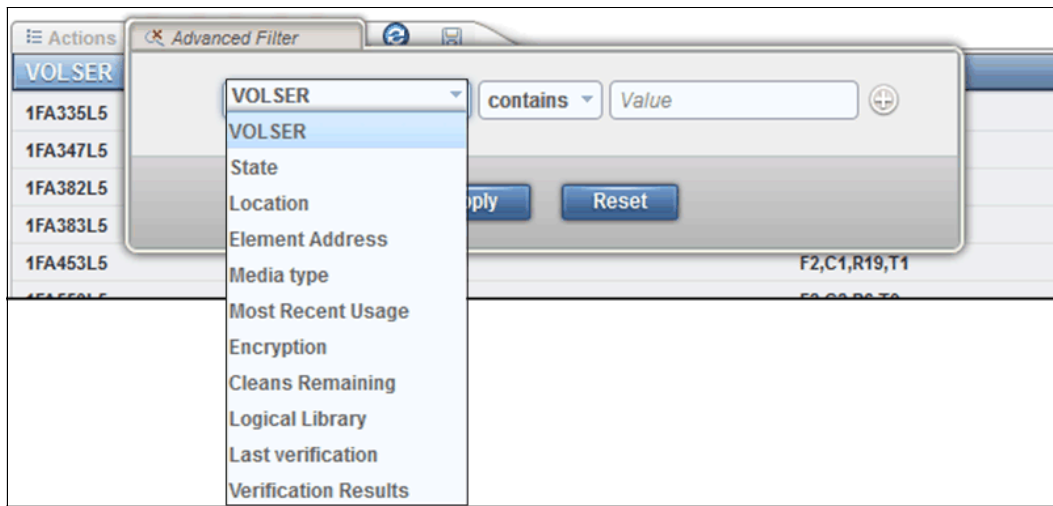


Figure 5-78 Advanced filter options for cartridges

The Cartridge database can be saved to a .csvfile 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), then select **Export Table Data**, as shown in Figure 5-79.




Figure 5-79 Export list to file

Viewing cleaning cartridge status

IBM supplies a cleaning cartridge with the 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 library, sort the column by Cleans Remaining by using the arrow, as shown in Figure 5-80.

The screenshot shows a table with columns: VOLSER, State, Location, and Cleans Remaining. The 'Cleans Remaining' column is highlighted in yellow and has a downward arrow indicating it is sorted. The table data is as follows:

VOLSER	State	Location	Cleans Remaining
CLNU57L1	✓ Normal	F1, C4, R44, T1	35
Z758NLT	✓ Normal	F1, C2, R2, T6	
Z669NLT	✓ Normal	F1, C1, R19, T6	
Z386NLT	✓ Normal	F1, C2, R3, T9	
Z186NLT	✓ Normal	F1, C3, R1, T5	
Z063NLT	✓ Normal	F1, C2, R1, T2	
Y734NLT	✓ Normal	F1, C1, R2, T1	
Y657NLT	✓ Normal	F1, C3, R1, T4	
Y652NLT	✓ Normal	F1, C1, R1, T5	
Y617NLT	✓ Normal	F1, C3, R1, T7	
Y303NLT	✓ Normal	F1, C3, R1, T2	
X920NLT	✓ Normal	F1, C4, R1, T1	
X01492L4	✓ Normal	F1, C3, R30, T4	

Figure 5-80 Number of cleanings remaining

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 128.

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:

- ▶ Normal
Cartridge status is normal
- ▶ Importing
Cartridge import consists of physically inserting the cartridge into an I/O slot, moving the cartridge to a physical storage slot within the library, and having the cartridge assigned to a logical library. The import is typically complete when the host application performs a move of the cartridge. Alternatively, you can use the GUI/CLI to assign it to a logical library directly. This state is only used for data cartridges.
- ▶ Uncertain Barcode
There is a mismatch between a cartridge barcode label scanned and the copy of the barcode stored on cartridge memory. The library attempts to resolve the issue on the next mount/demount cycle of this cartridge. Contact IBM for support if this persists.
- ▶ Assignment Required
The cartridge is in a slot but not assigned to a logical library because an empty import or export element address is not currently available for assignment. The cartridge must be manually assigned to the logical library or removed from the library.
- ▶ Export Queued
Queued to be moved to the service Magazine.
- ▶ Failed Verification
You might need to migrate the data to another cartridge.

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.

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 on the 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 managing cartridges and assigning them 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/Reassign**, as shown in Figure 5-81 on page 107.

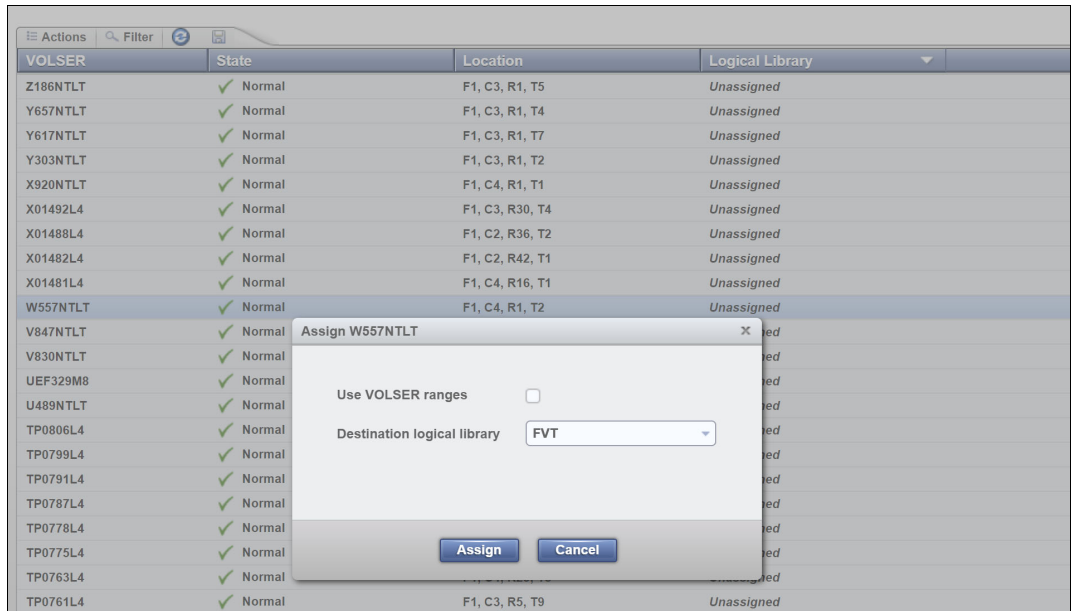


Figure 5-81 Assign volume example

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**. A pop-up window opens with a warning message that states this action might result in synchronization problems with the application (Figure 5-82). To avoid this scenario, the preferred method is to use the application to eject cartridges. If you want to continue, then select **Yes**.

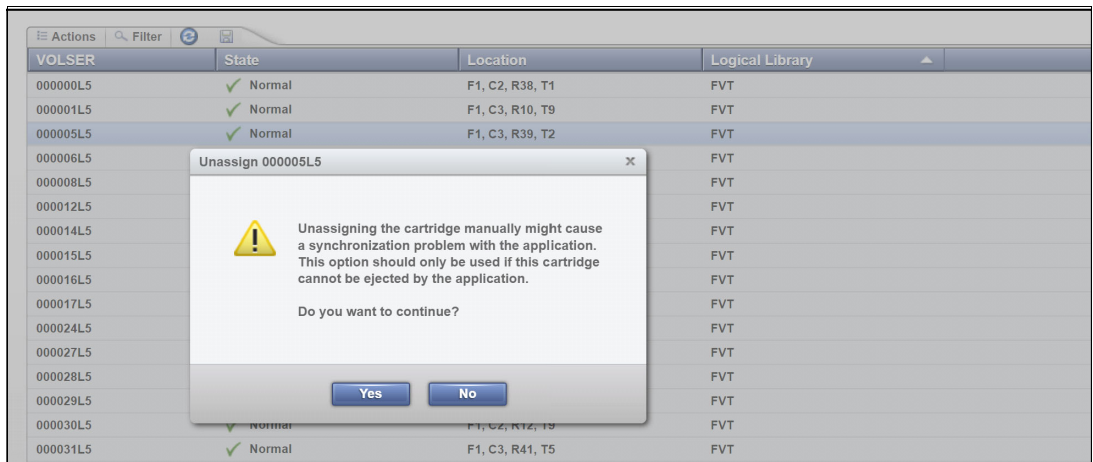


Figure 5-82 Unassign a volume example

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 5-83.

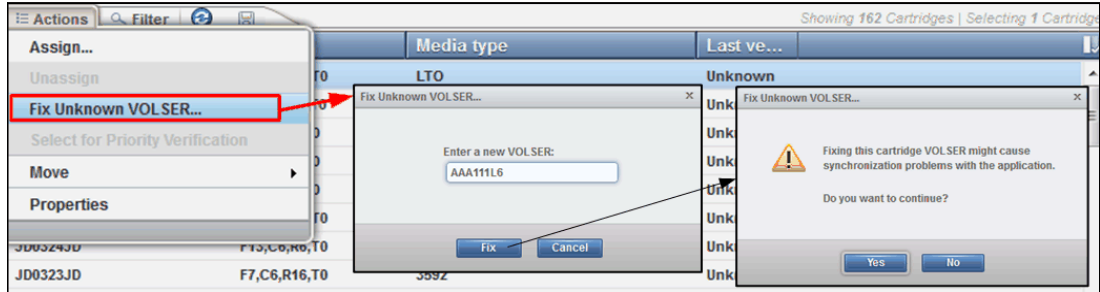


Figure 5-83 Fix unknown volser example

Table 5-3 lists the possible cartridge type values.

Table 5-3

Media	Cartridge Type
LTO media	L9,L8,LY,LZ

VOLSER

Note: This action might result in synchronization problems with the application. An audit type of action for your application might be required.

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 5-84.

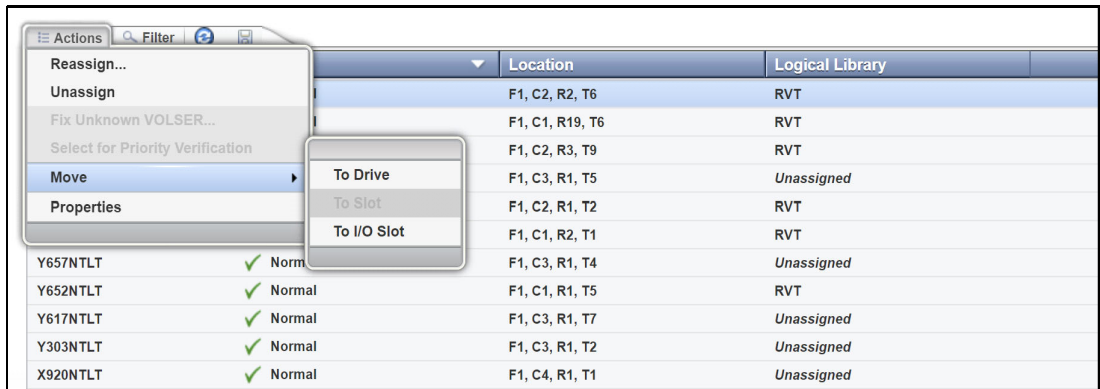


Figure 5-84 Move cartridge menu

Select **Actions Move To Drive**, **Actions Move To Slot**, or **Actions Move To I/O Slot**, depending on the location to which the cartridge is being moved.

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 (Figure 5-79 on page 105) 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.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1	Mount History						Mount	>>>	>>>	>>>	>>>	>>>	Life	>>>	>>>	>>>
2							Media	>>>	Host	>>>	Drive	Tape Alert	Mounts	WRetries	WPerms	RRetries
3							Write	Read	Write	Read	Residency	Media	Media	Media	Media	Media
4							MBs	MBs	MBs	MBs	Minutes	_1-64	Count	Count	Count	Count
5	Date and Time	VolSer	Frame	Drive	LogLib	EAddr	MBs	MBs	MBs	MBs	Minutes	Media	Media	Media	Media	Media
6	2023 05 18 15:35	_CLNU57L1	1	1	_	263	0	0	0	0	10	0	15	0	0	0
7	2023 04 21 11:06	_F32152L9	1	5	_FVT	257	205108	77223	204801	77106	20	0	1084	37	0	0
8	2023 04 21 11:26	_F32162L9	1	2	_FVT	258	102554	102554	102599	102401	21	0	1363	57	0	0
9	2023 04 21 11:35	_F32140L9	1	10	_FVT	262	102554	102554	102401	102401	24	0	1117	27	0	0
10	2023 04 21 11:47	_F32162L9	1	1	_FVT	263	102554	102554	102401	102401	41	0	1364	57	0	0
11	2023 04 21 11:59	_F32140L9	1	7	_FVT	260	102554	102554	102401	102401	12	1.21E+09	1118	27	0	0
12	2023 04 21 12:10	_F32140L9	1	5	_FVT	257	102554	102554	102599	102401	12	0	1119	27	0	0
13	2023 04 21 12:28	_F32142L9	1	10	_FVT	262	102561	102561	102401	102401	9	0	1313	24	0	0
14	2023 04 21 12:40	_F32177L9	1	2	_FVT	258	102554	102554	102401	102401	5529	0	516	26	0	0
15	2023 04 25 08:39	_1FC042L6	1	6	_FVT	259	0	0	0	0	1	0	178	1	0	0
16	2023 04 25 08:39	_1FC042L6	1	10	_FVT	262	0	0	0	0	1	0	178	1	0	0
17	2023 04 25 08:40	_2FA406L5	1	1	_FVT	263	0	0	0	0	3	0	153	12402	0	0
18	2023 04 25 08:41	_1FC006L6	1	6	_FVT	259	0	0	0	0	1	0	3725	9	0	0
19	2023 04 25 08:41	_1FC006L6	1	9	_FVT	261	0	0	0	0	2	0	3725	9	0	0
20	2023 04 25 08:42	_2FA369L5	1	5	_FVT	257	0	0	0	0	1	0	2	6	0	0
21	2023 04 25 08:43	_1FC006L6	1	1	_FVT	263	0	0	0	0	6	0	3725	9	0	0
22	2023 04 25 11:00	_F32177L9	1	5	_FVT	257	102554	102554	102401	102401	35	0	517	26	0	0
23	2023 04 25 11:09	_F20058L9	1	1	_FVT	263	102554	102554	102599	102401	12	0	530	7	0	0
24	2023 04 25 11:22	_F32152L9	1	7	_FVT	260	102554	102554	102401	102401	33	0	1085	37	0	0
25	2023 04 25 11:34	_F32177L9	1	9	_FVT	261	102554	102554	102599	102401	9	0	518	26	0	0
26	2023 04 25 11:46	_F32177L9	1	10	_FVT	262	102554	102555	102401	102401	194	0	519	26	0	0
27	2023 04 25 12:01	_F32152L9	1	7	_FVT	260	102554	102554	102599	102401	51	0	1086	37	0	0
28	2023 04 25 12:10	_F32178L9	1	2	_FVT	258	102554	102554	102401	102401	12	0	891	22	0	0
29	2023 04 25 12:22	_F32178L9	1	5	_FVT	257	102554	102554	102401	102401	12	0	892	22	0	0
30	2023 04 25 12:34	_F32178L9	1	6	_FVT	259	102554	102554	102401	102401	217	0	893	22	0	0

Figure 5-85 Exported mount history example

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.
- ▶ **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 in the ready mode. Ready mode means when data can be moved but it is not. 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 of 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 in the ready mode. Ready mode means 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 which includes setup and overhead. 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 because of recording error recovery and media defects for example.
- ▶ 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. If the media contains encrypted data, the value is 1. If the does not contain encrypted data, the value is 0. If the drive is unable to determine whether the media contains encrypted data, the value is a blank space.
- ▶ Crypto Rekey. Whether a cartridge was rekeyed. If a cartridge was rekeyed during the last mount, the value is 1. If a cartridge was not rekeyed during the last mount, the value is 0.
- ▶ Mount Host Write. The number of megabytes (MB) that was written during a mount.
- ▶ Mount Host Read. The number of 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.

5.5.2 Cartridges by Logical Library page

Use the Cartridges by Logical Library page to create and manage logical libraries. From this page, you can also assign, unassign, or move cartridges between the logical libraries.

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 5-86 on page 112 shows the main window for Cartridges by Logical Library and the available selections for the default columns that you want to view.

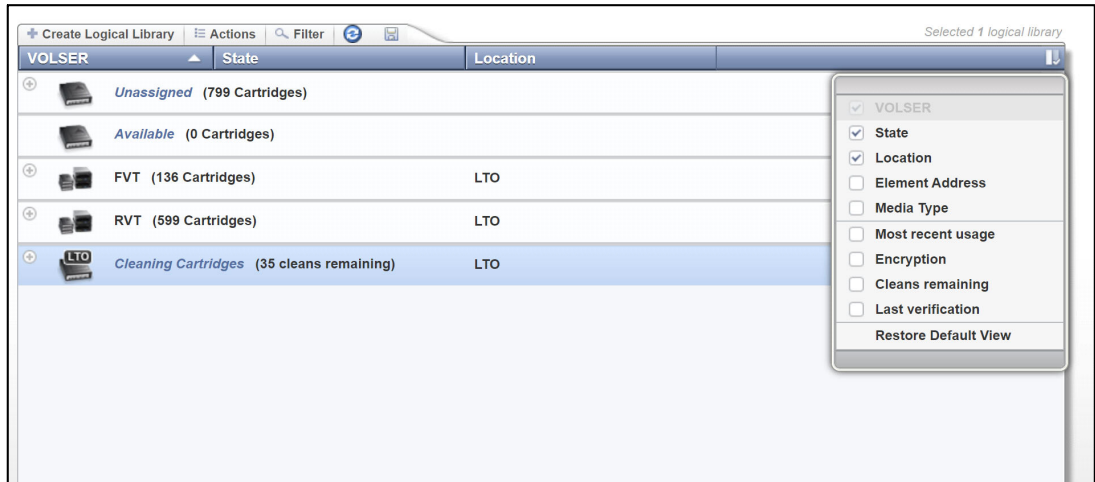


Figure 5-86 Cartridges by logical library page

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 each 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. The number of cleanings remaining on a specific cleaning cartridge can also be viewed in the Cleans Remaining column on the Cartridges page, as shown in Figure 5-87 on page 113.

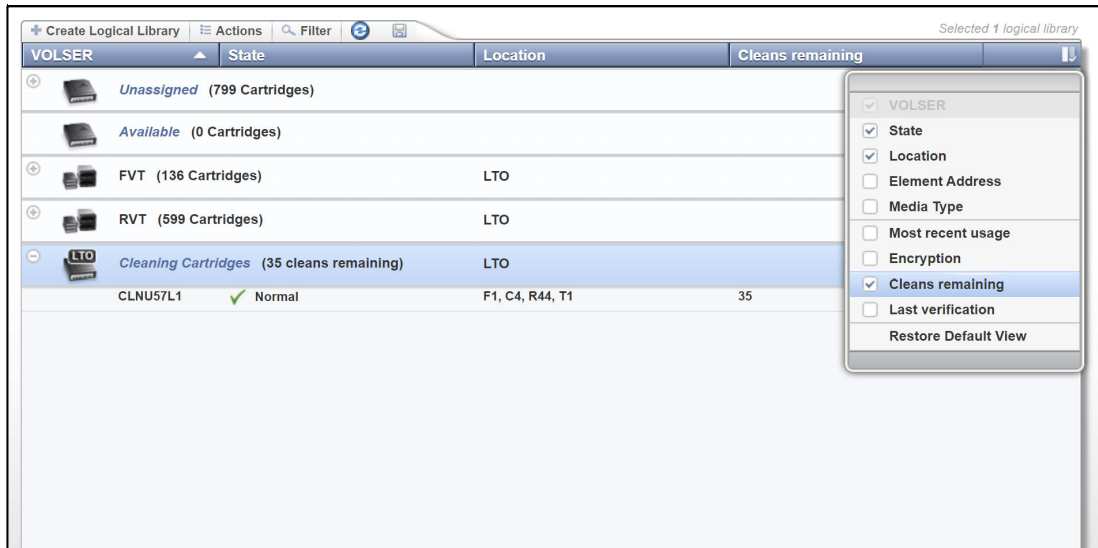


Figure 5-87 Cleaning cartridge information

You can optionally automatically eject expired cleaning cartridges. To enable this feature, select **Settings Library Advanced**, and select **Automatically**, as shown in “Advanced” on page 128.

Note: Cleaning cartridges are shared among all logical libraries.

Assigning or unassigning cartridges

A bar code label with a 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.

To change which cartridges are assigned to a logical library, modify 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 select **Actions Assign** (Figure 5-88 on page 114).

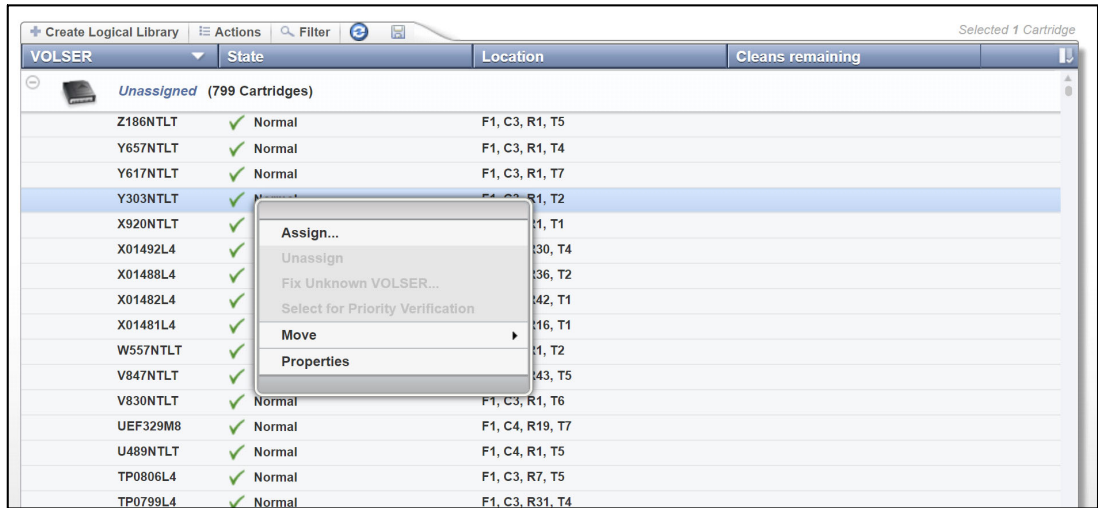


Figure 5-88 Assign cartridges to a logical library

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.

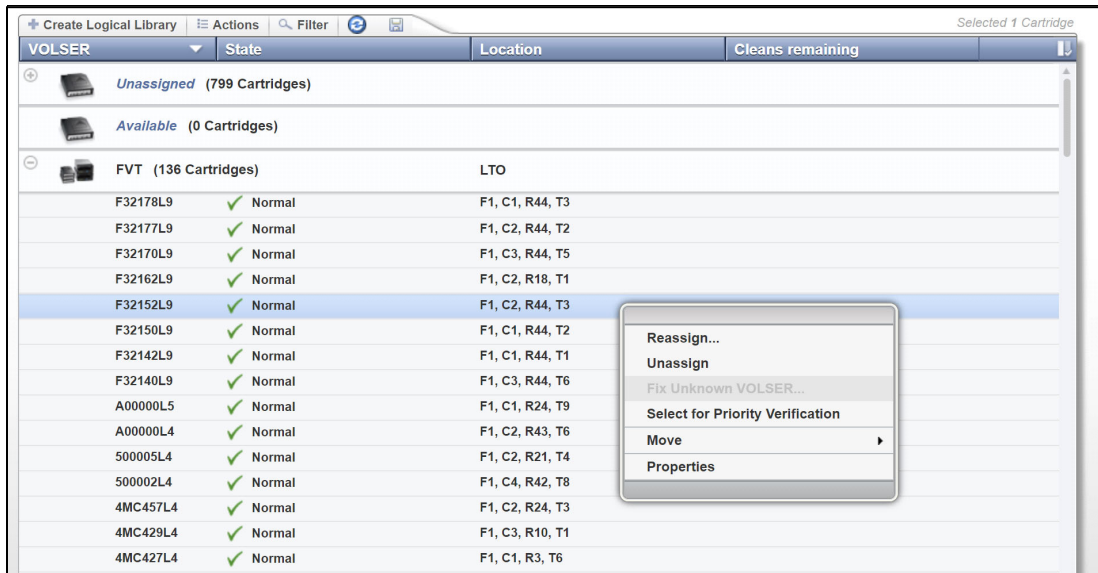


Figure 5-89 Unassigning cartridges manually

Depending on the preferred cartridge location, 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

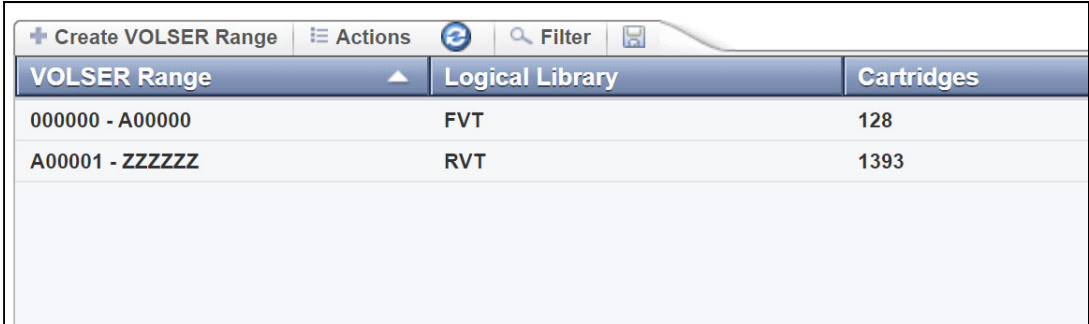
For more information about creating and managing logical libraries, including setting, or modifying the method of encryption that is used by logical libraries, see 5.3.1, “Creating a logical library” on page 81.

5.5.3 VOLSER ranges

VOLSER ranges are used to assign cartridges to specific logical libraries. 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



VOLSER Range	Logical Library	Cartridges
000000 - A00000	FVT	128
A00001 - ZZZZZZ	RVT	1393

Figure 5-90 Create VOLSER range page

When you insert a cartridge into the library and when 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, a logical library of LTO drives is created and named LogicalLibrary1 for VOLSERs that range from ABC000 to ABC999. If a cartridge with VOLSER ABC123 is inserted into the library, the library recognizes that the VOLSER belongs in the range and assigns it to LogicalLibrary1.

Inserting a cartridge that is outside of any defined VOLSER ranges makes it 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 Diamondback 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** (Figure 5-91 on page 116).

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.

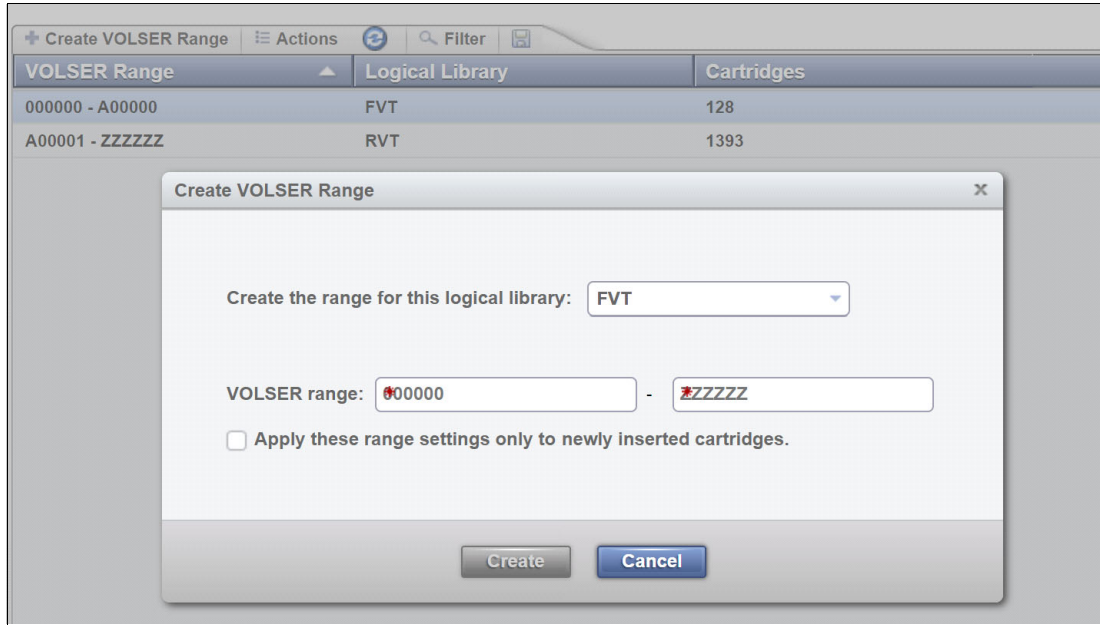


Figure 5-91 Create or modify VOLSER range

Reassigning a VOLSER range

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 5-92. 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.

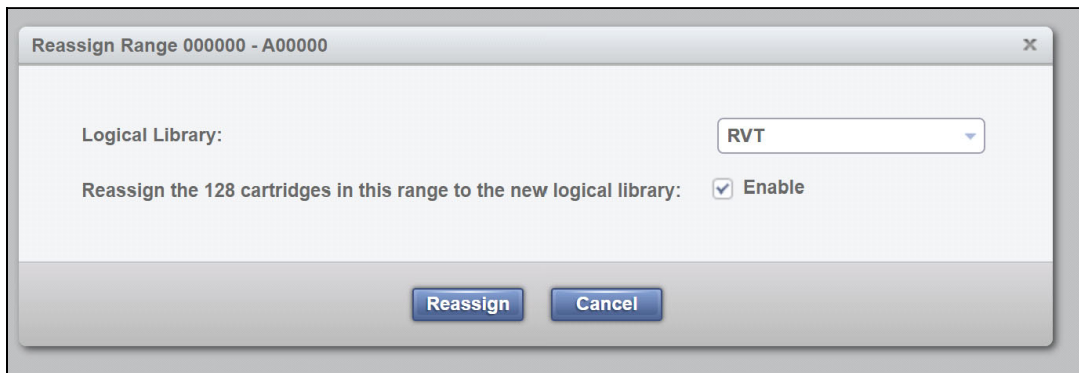


Figure 5-92 Reassign VOLSER range

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.

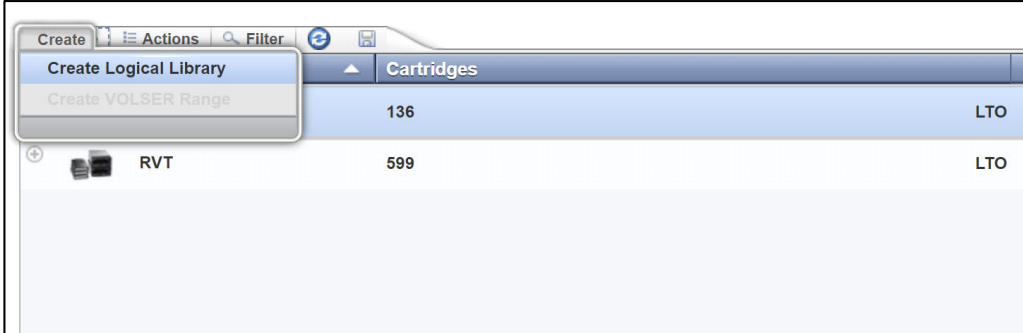
Note: Any action that causes cartridge assignments to change can result in application synchronization problems. To avoid this scenario, first use the application to eject any cartridges that are assigned to the selected VOLSER range.

5.5.4 VOLSER range by logical library

VOLSER ranges are used to automatically assign cartridges to specific logical libraries. Ranges are created when creating your logical libraries or assigned to the library later. VOLSER ranges can be modified or reassigned on the Logical Library page. The following information is presented for using VOLSER ranges:

- ▶ Creating and managing logical libraries
- ▶ Creating, modifying, or deleting VOLSER ranges

Figure 5-93 shows the VOLSER Range by Logical Library window.



The screenshot shows a software window titled "VOLSER Range by Logical Library". It features a menu bar with "Create", "Actions", and "Filter". A dropdown menu is open under "Create", showing "Create Logical Library" and "Create VOLSER Range". Below the menu is a table with the following data:

	Cartridges	
	136	LTO
+ RVT	599	LTO

Figure 5-93 VOLSER range by logical library

Using VOLSER ranges

When you insert a cartridge into the library and when 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, a logical library of LTO drives is created and named LogicalLibrary1 for VOLSERs that range from ABC000 to ABC999. If a cartridge with VOLSER ABC123 is inserted into the library, the library recognizes that the VOLSER belongs in the range and assigns it to LogicalLibrary1.

Inserting a cartridge that is outside of any defined VOLSER ranges makes it 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 Diamondback 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

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 5-92 on page 116. 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.

5.6 Access menu

Use the Diamondback management GUI pages, which are under the Access icon, as shown in Figure 5-94, to view, create, and assign users and their roles.



Figure 5-94 Access menu options

Note: The options and actions under the Access icon are available only to users with the administrator role.

5.6.1 Users

From the Users page administrators can create and manage users, map users to a role, view which users have active sessions and how many connections they have.

Note: With the Diamondback tape library, password protection is always enabled and all users are required to sign in with a username and password.

Selecting **Users** from the Access menu gives an overview of all configured users, as shown in Figure 5-95 on page 119.




Figure 5-95 Users tab contents

Table 5-4 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 5-95.

Table 5-4 User page components

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	Last login
9	User role
10	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 5-5 lists the available actions.

Table 5-5 Local and remote authentication actions

Actions	Local authentication	Remote authentication
Create User	x	
Map to Role	x	
Reset Password	x	
Unlock	x	
Delete	x	
Disconnect	x	x
Connections	x	x
Properties	x	x

Local authentication

With local authentication, each Diamondback tape library maintains a separate database of usernames 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 Diamondback 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 of groups. If one of these groups matches a customized role within the Diamondback, the user is granted the permissions that are associated with the role on the library.

Note: The name of the group on the remote authentication server must be the same as the name of the role on the Diamondback tape library for the group to be matched to that role.

The library supports 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 240.

For more information about enabling remote authentication, see “Enable remote authentication” on page 137 and 5.8.1, “Remote support by using an external TSSC” on page 154.

Creating users

You can create a maximum of 80 users, each with a unique username and password.

Each user account is mapped to a role that defines the pages, which the user can view and 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 5.6.2, “Roles” on page 125.

To create a user, select **Create User** in the Users window. The Create User window opens, as shown in Figure 5-96 on page 121. Enter the username and select a role for the user. The user receives the role that is chosen from the Role drop-down list. Ensure that any roles beyond the default are configured before creating the user.

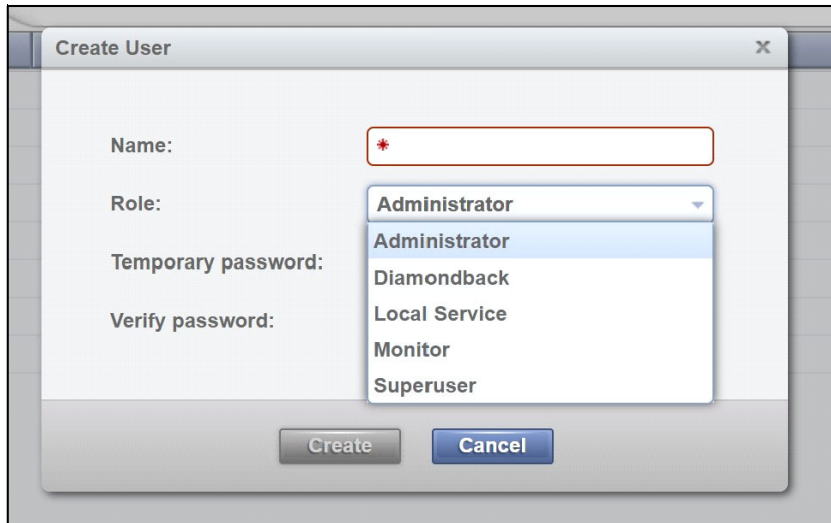


Figure 5-96 Create user window

For more information about creating and managing roles, see 5.6.2, “Roles” on page 125. A user can receive the temporary password from an administrator that 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 144.

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 (Figure 5-97).




Figure 5-97 User created confirmation message

Note: It is a best practice 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 (Figure 5-98 on page 122).

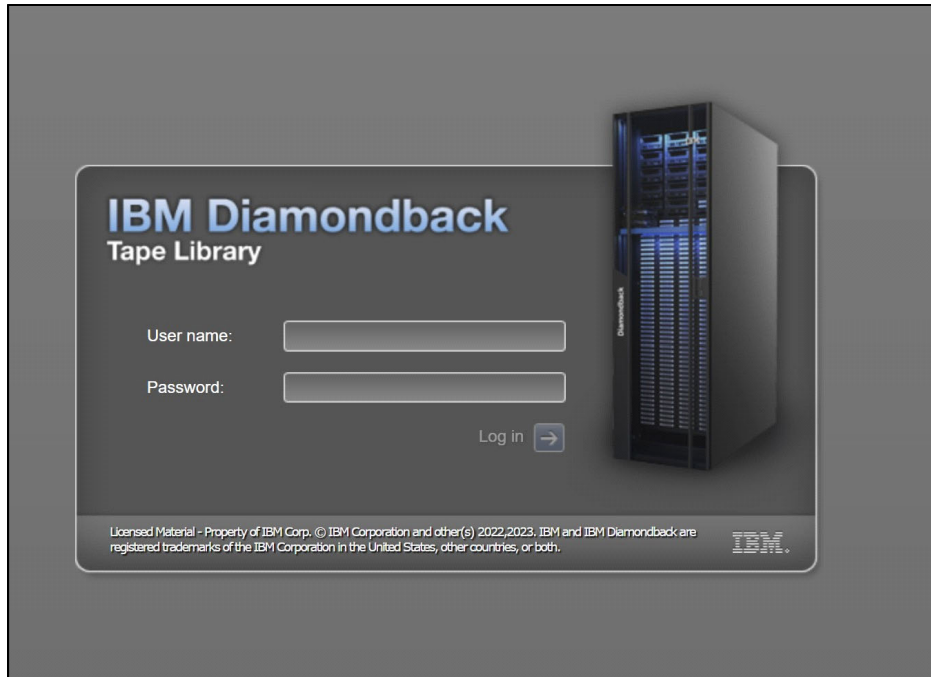


Figure 5-98 User login screen

The user is then presented with the change password window (Figure 5-99).



Figure 5-99 Change password prompt

The new password must conform to the rules defined in the Password and Session Policy. For more information, see “Password and Session Policy” on page 144.

Note: The username and password are case-sensitive.

After the user sets their password, the user can log in with the new password in the initial login window, as shown in Figure 5-98 on page 122. After the user logs in with the new password, the user can access the Diamondback 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 by using the Actions menu (Figure 5-100).

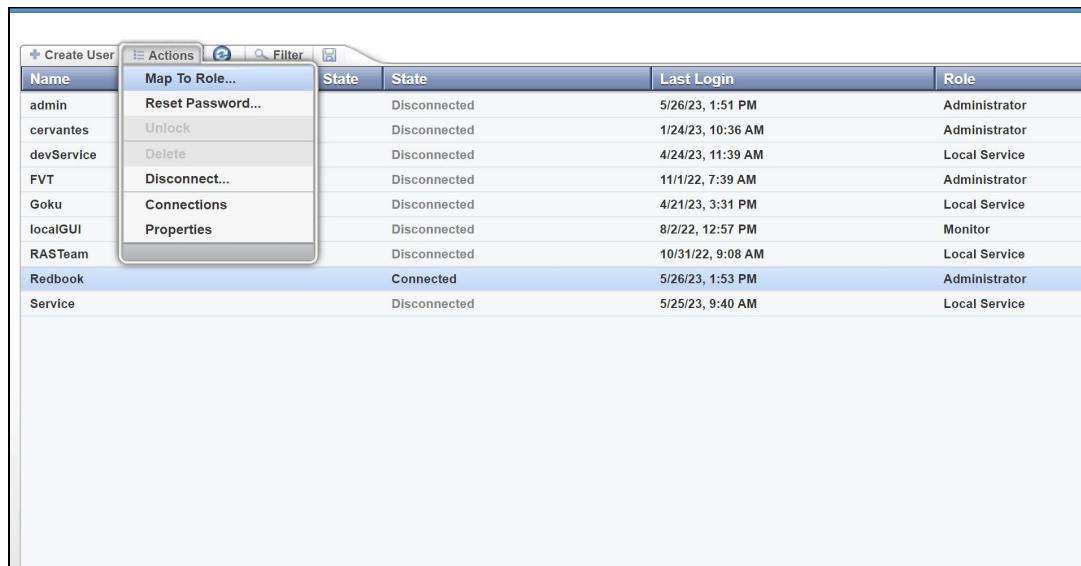


Figure 5-100 User Actions menu

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.
- ▶ **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 121.
- ▶ **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 5-101 on page 124.

Name	Locked State	State	Last Login	Role
admin		Connected	5/30/23, 8:32 AM	Administrator
cervantes	🔒	Disconnected	1/24/23, 10:36 AM	Administrator
devService		Disconnected	4/24/23, 11:39 AM	Local Service
FVT		Disconnected	11/1/22, 7:39 AM	Administrator
Goku		Disconnected	4/21/23, 3:31 PM	Local Service
localGUI		Disconnected	8/2/22, 12:57 PM	Monitor
RASTeam		Disconnected	10/31/22, 9:08 AM	Local Service
Redbook		Disconnected	5/26/23, 1:53 PM	Administrator
Service		Disconnected	5/25/23, 9:40 AM	Local Service

Figure 5-101 Locked user example

If the user lost their password, any administrator can reset the password. If the user is locked, an administrator must unlock the user. An account must be unlocked to change its password.

- ▶ **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 Diamondback.
- ▶ **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 5-102.

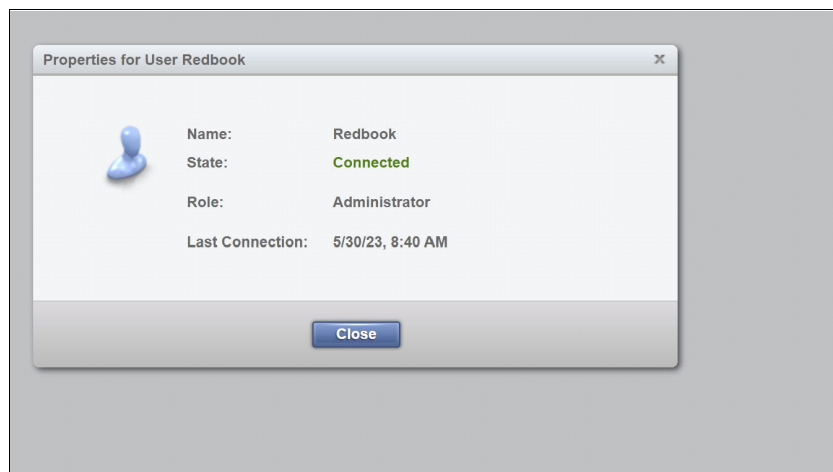


Figure 5-102 User properties

Locked out

If you are locked out of the Diamondback management GUI and unable to have your account password reset, you can reset the default administrator account username and password. Use the Access Recovery button located on the back panel of the Diamondback tape library. This button resets the username and password to the default logon settings for 15 minutes.

Use this reset button only for emergencies. For example, use it if the remote authentication server is down or if the administrator loses or forgets their password and no other administrators are able to access the system.




Figure 5-103 Library back panel

Perform the following steps to reset the default username and password:

1. Press the Access Recovery button (1 in Figure 5-103) on the front of the electronics tray. If a remote connection is not available, connect a laptop to the service port (2 in Figure 5-103).
2. Log in to the GUI by using **admin** for the username and password.
3. Go to your username on the upper right of the toolbar and select **Modify Password**.
4. Enter your current password and your new password and click **Modify**.

5.6.2 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 5-104 shows the Role menu.




Figure 5-104 Role menu example

Preset roles

The Diamondback tape library has four preset roles, each with a different level of access to the Diamondback management GUI. It is also possible to create 16 other 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.
- ▶ **Local Service.** Users that are mapped to this role can view all pages that are available to a monitor role. Also, 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 Diamondback 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.

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 (Figure 5-105). All roles can access all logical libraries in the Diamondback tape library.

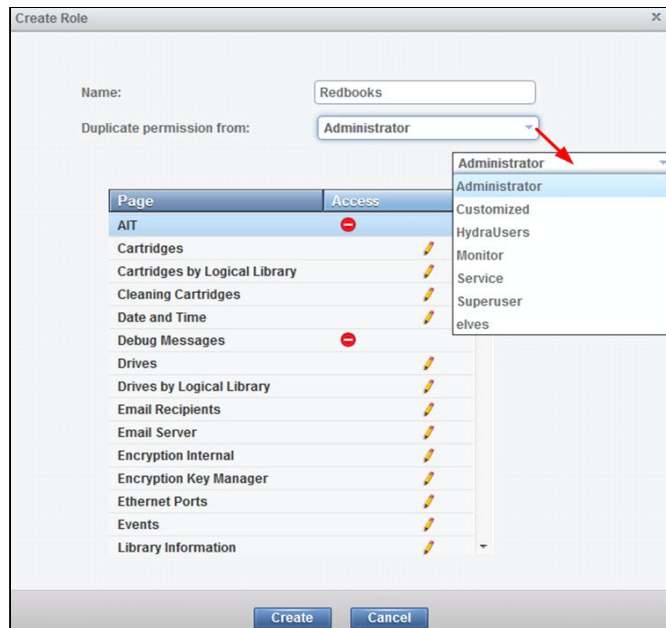


Figure 5-105 Create role window and options

To remove a role, select the role. Then, click **Access Roles** → **Actions** → **Delete**. You can delete a role only when there are no users are assigned to that role.

5.7 Settings

Use the Settings menu as shown in Figure 5-106 to configure overall library options and settings. Each page is described in the following sections.

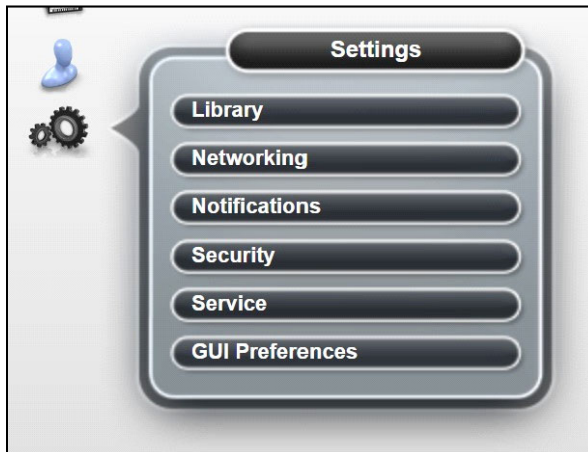


Figure 5-106 Settings menu

5.7.1 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 (Figure 5-107) or to synchronize with the NTP server (Figure 5-108 on page 128). If you select to use the NTP server, you can optionally insert a primary and secondary NTP server and test the connection from this page.

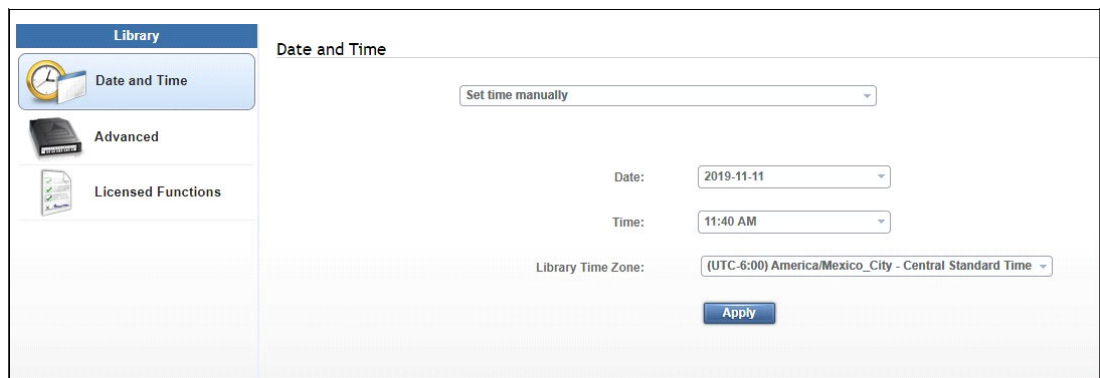


Figure 5-107 'Set time manually' options

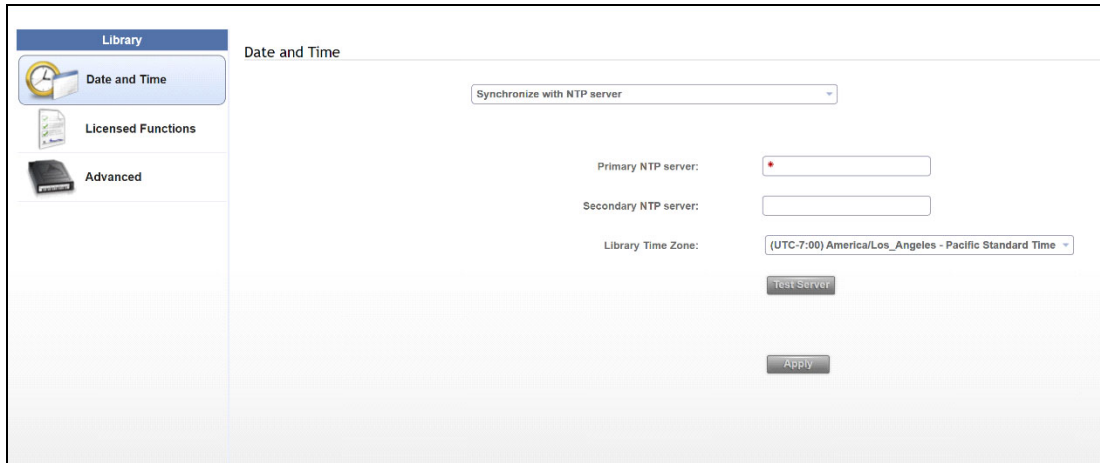


Figure 5-108 Synchronize with NTP server options

Advanced

Use the Advanced page to manage expired cleaning cartridges, and REST over SCSI as shown in Figure 5-109.

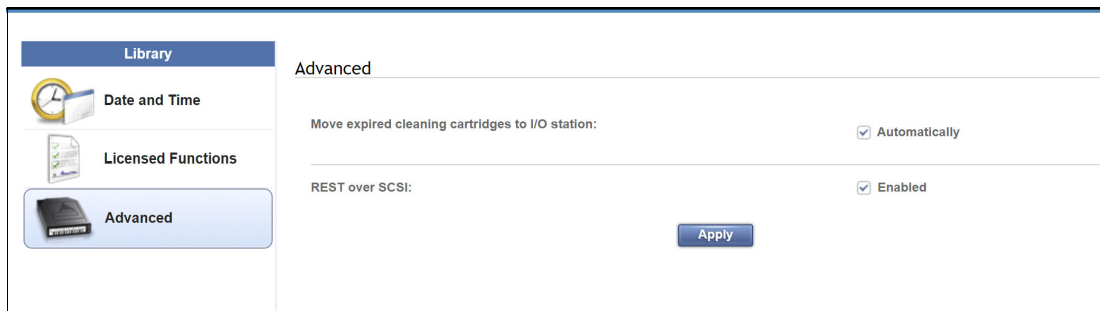


Figure 5-109 Advanced options with single accessor

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 Diamondback 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 for each cartridge on a memory chip inside the cleaning cartridge. This prevents an expired cleaning cartridge from being reused.

REST over SCSI

Use this option to enable, which is the default, or disable REST over SCSI. REST over SCSI is described in detail in Chapter 6, “IBM Diamondback REST API” on page 57.

Licensed Functions

The following section describes the licensing of additional features to the library. These are extended library capabilities that become available when the license key is applied through the GUI:

- *Full Capacity* is the listed example. It increases storage from the initial capacity of 800 slots to the full capacity of 1584 slots.

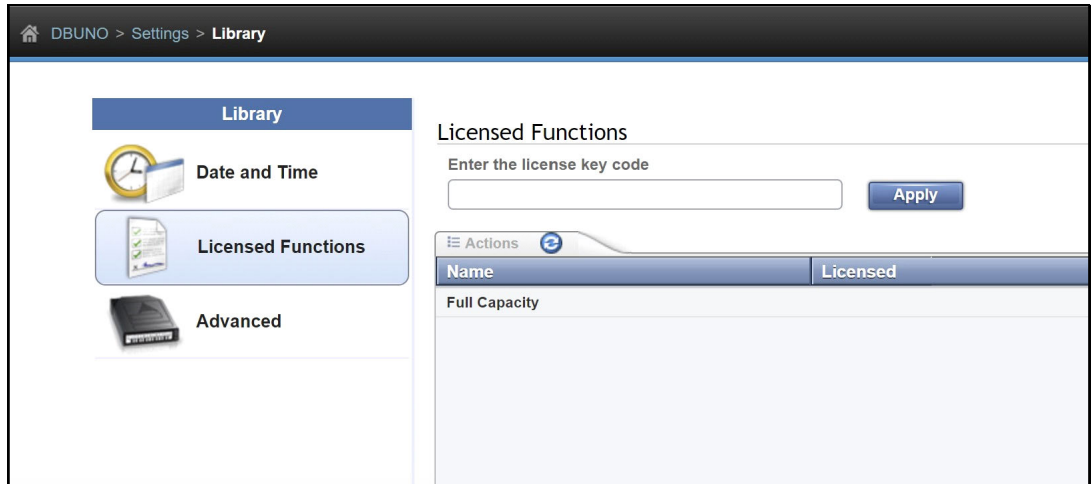


Figure 5-110 Licensed function page

It is possible to remove a licensed function by right-clicking a licensed function and selecting **Remove License Key** as shown in Figure 5-111.

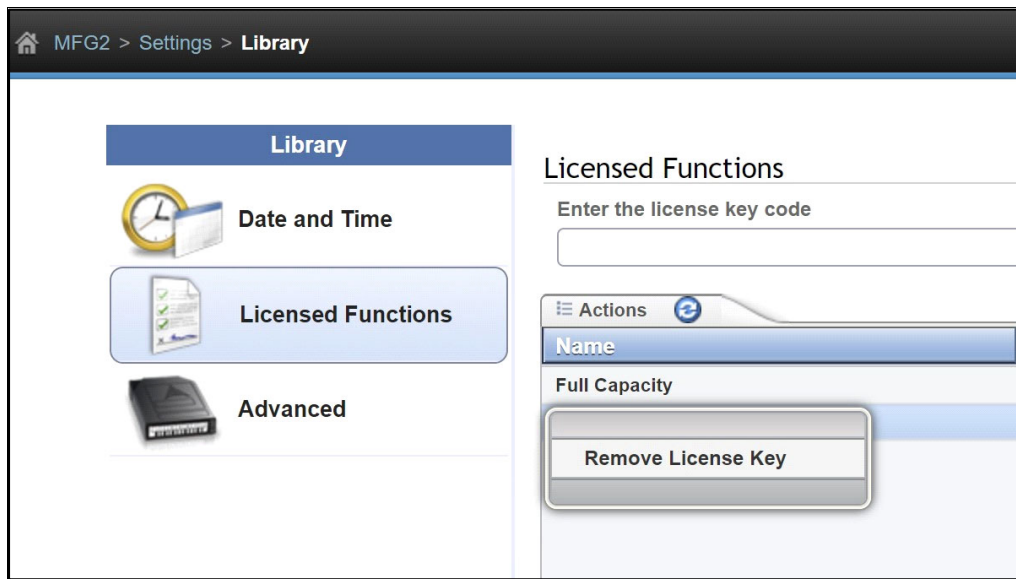


Figure 5-111 Remove Licensed function panel

Note: Remove License function for Full Capacity is currently disabled.

5.7.2 Networking

Select Networking on the Settings menu to access the Networking page to set up management Ethernet ports.

Management Ethernet Ports

Use the Management Ethernet Ports menu to display and configure the ports on the Diamondback Library. Each port can enable or disable a particular protocol.

By selecting the **Networking** option, the Ethernet Ports window opens, as shown in Figure 5-112, this displays the current settings of all ports and displays their status.

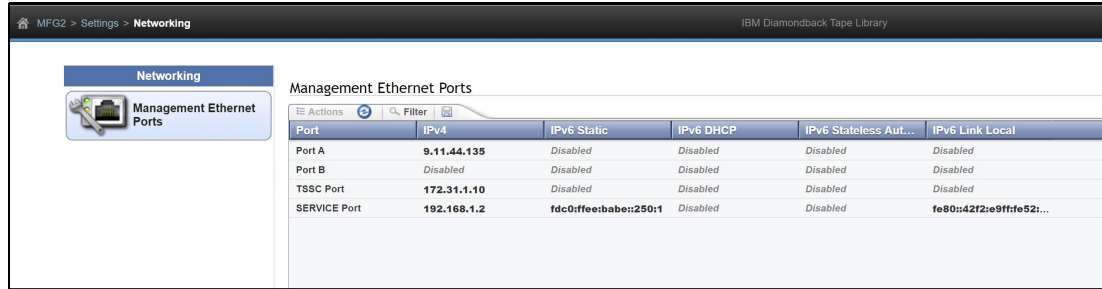


Figure 5-112 Ethernet ports in the Networking page

Use this window to change the Internet Protocol (IP) address that is used to access the management GUI. Individual 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 5-113.




Figure 5-113 Ethernet port settings menu

Modifying a port

To modify a port, highlight the port that you want to modify and select **Actions** → **Modify Port**. The Modify Port window opens, as shown in Figure 5-114 on page 131.

Use this window to change the IP address that is used to access the management GUI. Ports can be configured to use IPv4, IPv6, or both types of IP addresses.

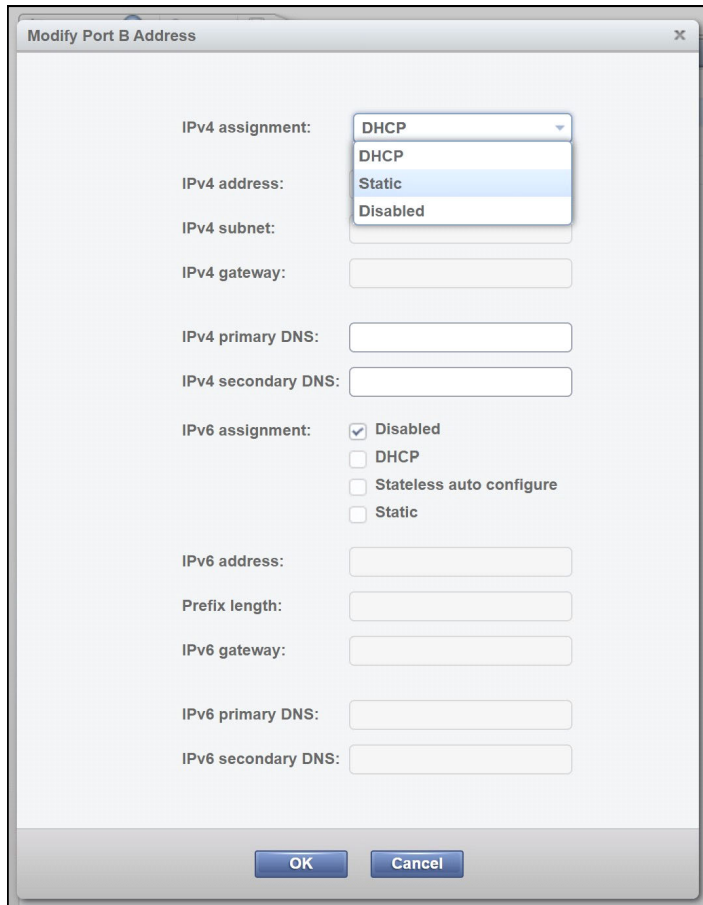


Figure 5-114 Modify Ethernet port settings

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. This option uses a combination of the router prefix, which identifies the subnet that is associated with a link and a host-generated interface identifier, which 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 5-115.

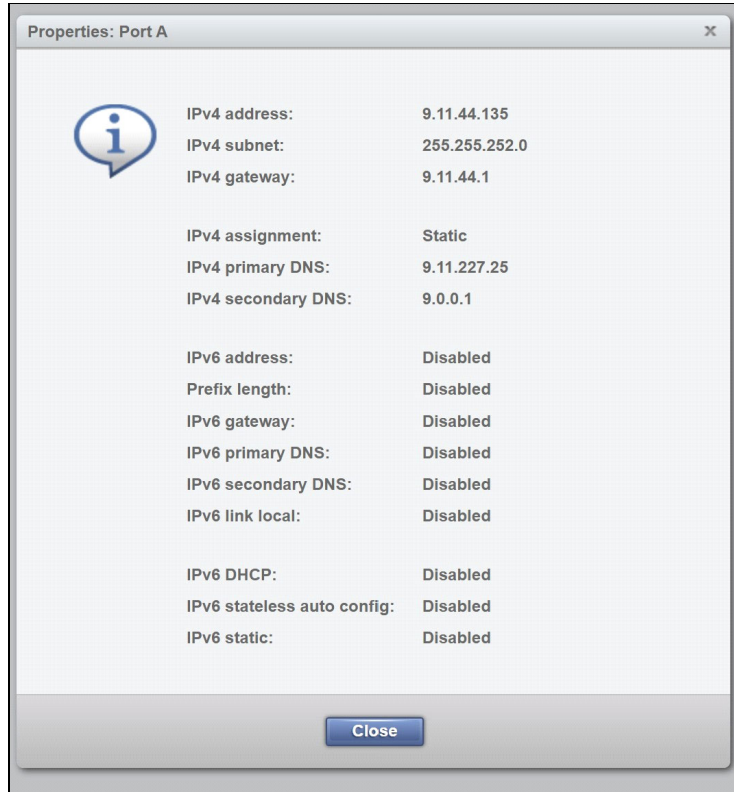


Figure 5-115 Ethernet port properties

5.7.3 Notifications

Select Notifications on the Settings menu to access the Notifications page to configure the sender information, recipient information, and library information for the Diamondback tape library notifications. The Notifications page contains several options, as shown in Figure 5-116 on page 133.

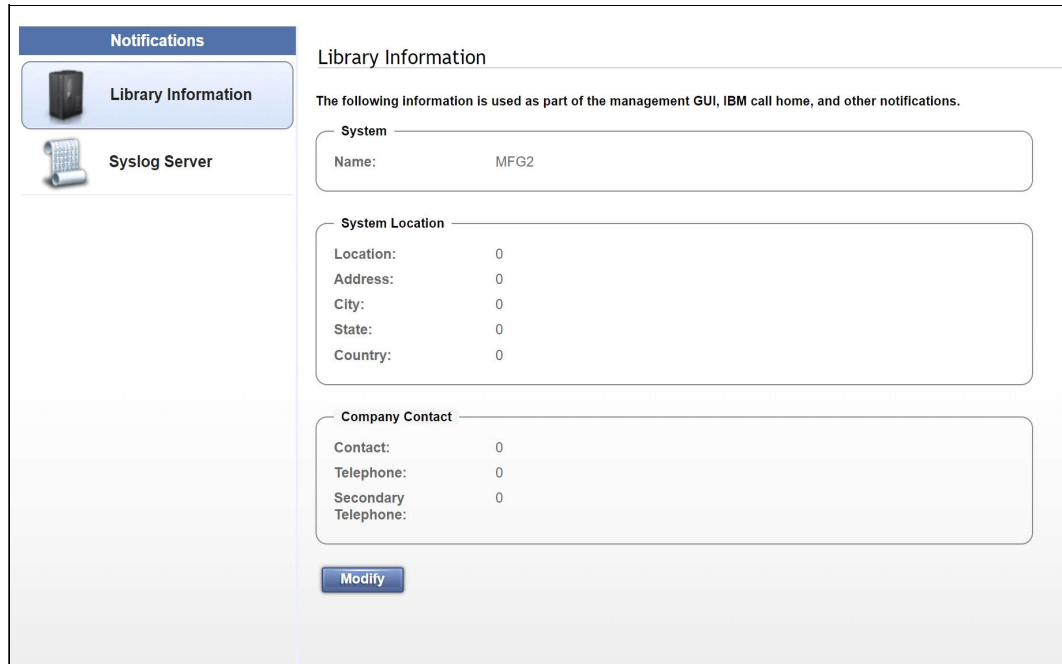


Figure 5-116 Notifications page

Library Information

The Library Information page displays the configured system name, system location, and company contact information for the Diamondback tape library. The library information is included in 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 5-116.

The Modify option opens the window that is shown in Figure 5-117 on page 134. On the Modify Library information window, you can set all of the system information and configure the library name, which can be changed concurrently.

The screenshot shows a dialog box titled "Modify Library Information". It contains several text input fields. The "System Name" field is filled with "Redbook". The other fields, "Location", "Address", "City", "State", "Country", "Contact", "Telephone", and "Secondary Telephone", are all filled with "0". At the bottom of the dialog, there are two buttons: "Modify" and "Cancel".

Figure 5-117 Modify library information window

2. Use this window to set the Diamondback 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.
3. After you complete all of the settings, click **Modify** to apply the changes.

System log (Syslog) server

Use the **Syslog Server** option on the Notifications menu to set up a destination syslog server, as shown in Figure 5-118 on page 135:

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**.

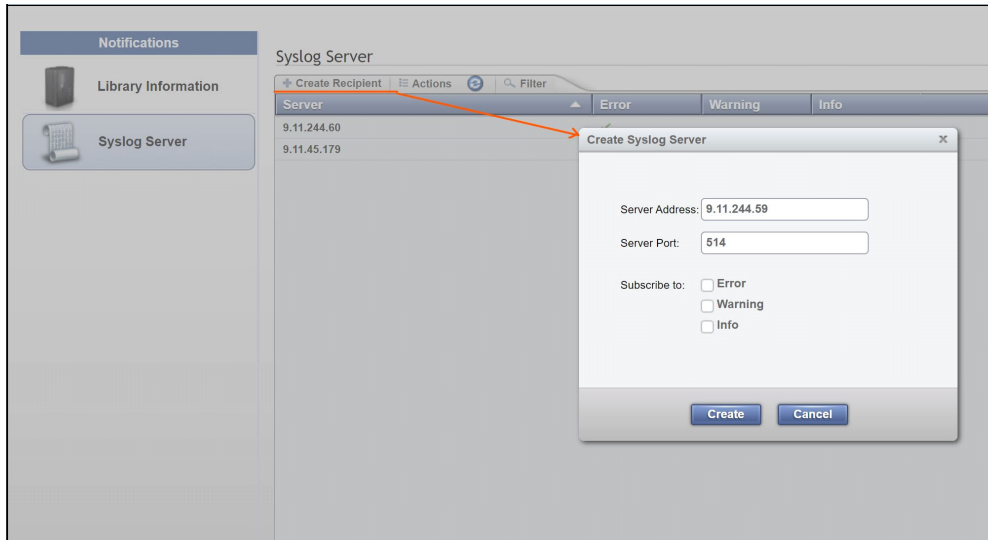


Figure 5-118 Syslog server panel

When system events occur, the Diamondback tape library creates a log of these events. You can configure the Diamondback 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 5-119. You can modify, delete, or send a test to the IP address of a syslog server.




Figure 5-119

Security

Select **Security** in the Settings menu to access the following library features:

- ▶ 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 Diamondback tape library. The default status for remote authentication is Disabled, as shown in Figure 5-120 on page 136.

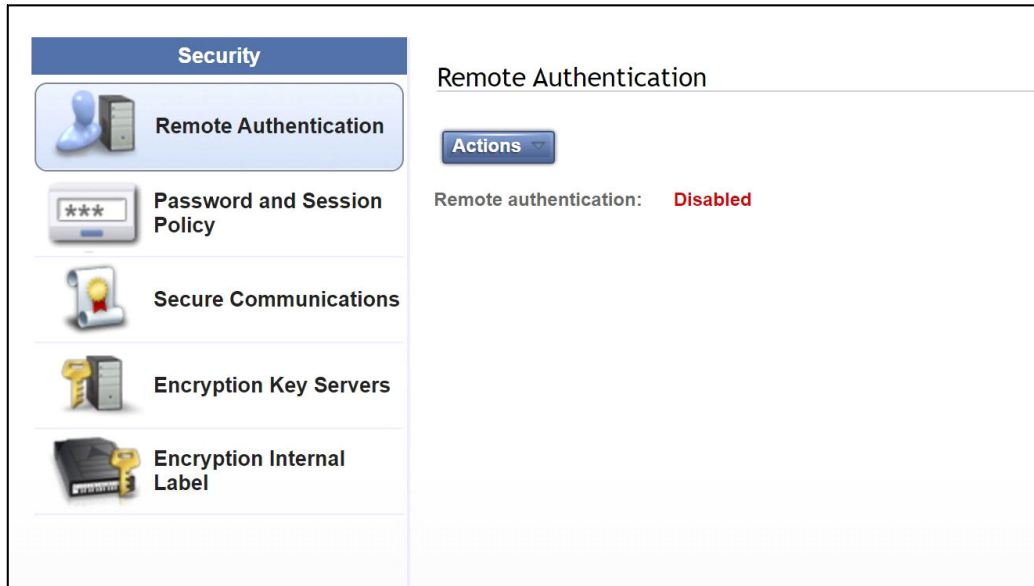


Figure 5-120 Security remote authentication page

LDAP is 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 228.

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 setup, see 5.6.1, “Users” on page 118.

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 Diamondback.

To enable remote authentication, the LDAP server has the following characteristics:

- ▶ A group is created and users are added to that group on your remote authentication server.
- ▶ On the Diamondback tape library, a custom role is created with a name that matches the group name on the remote authentication server. The Diamondback tape library supports 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 need 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 the provided scenario, users are mapped to groups and roles.

The concept is that a user is authenticated on the LDAP server and the Diamondback queries for all the groups in which the user belongs. These groups are mapped to customized roles within the Diamondback. 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 the Diamondback roles, the user inherits the access level of the first role match that is found in alphabetical order. For example, a user belongs to two groups that have corresponding Diamondback custom roles, `tech_admin` with Administrator access and `tech_service` with Service access. Therefore, the user 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 Diamondback 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 5.6.2, "Roles" on page 125.

Enable remote authentication

After you complete all required preparation and configuration, select **Actions** to enable remote authentication, as shown in Figure 5-120 on page 136.

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 5-121 on page 138.

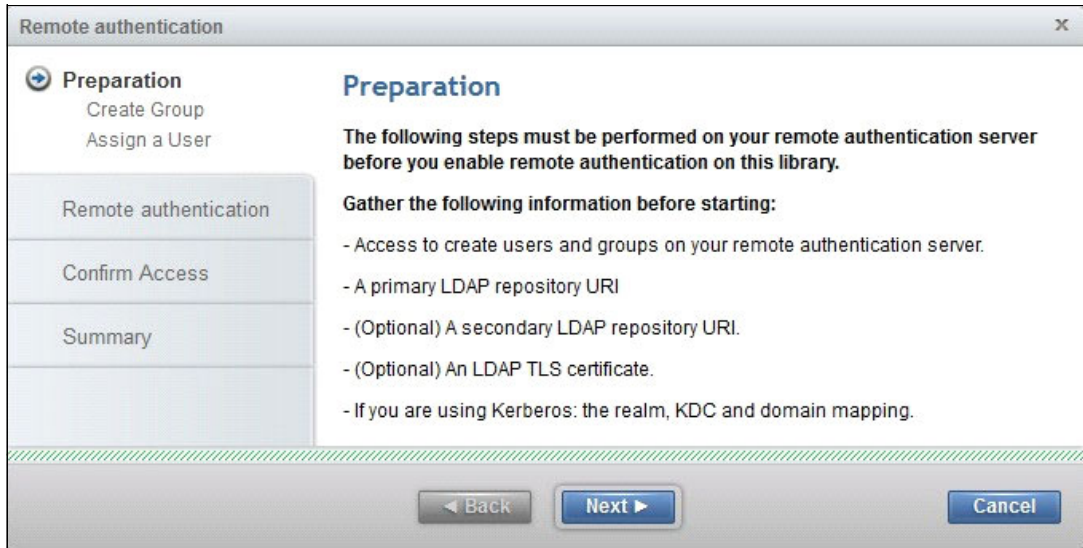


Figure 5-121 Remote authentication preparation window

The next two steps require you to confirm the custom Roles on the Diamondback 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 use Kerberos, select **Kerberos**. Otherwise, select **LDAP** (see Figure 5-122).

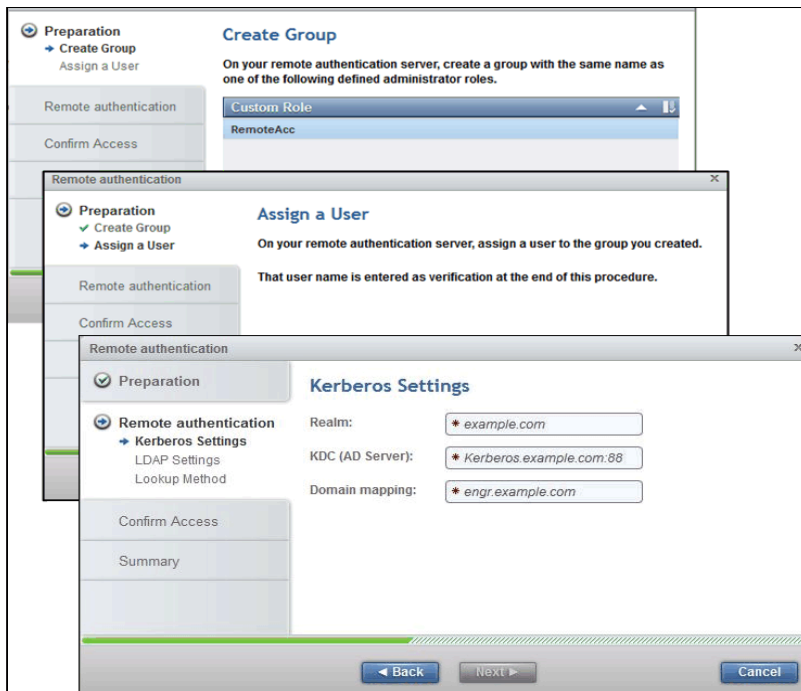


Figure 5-122 Remote authentication preparation

Kerberos

If you selected the Kerberos authentication method, the Kerberos Settings window opens (see Figure 5-123).



Figure 5-123 Kerberos settings window

Complete the following steps in the Kerberos Settings window:

- ▶ Enter the name of the realm setup for this machine.
- ▶ Enter the KDC (AD) server name.
- ▶ Enter the Domain mapping.

Figure 5-124 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 checkmark appears if a connection is made. If a connection cannot be made to the server, a red cross appears.

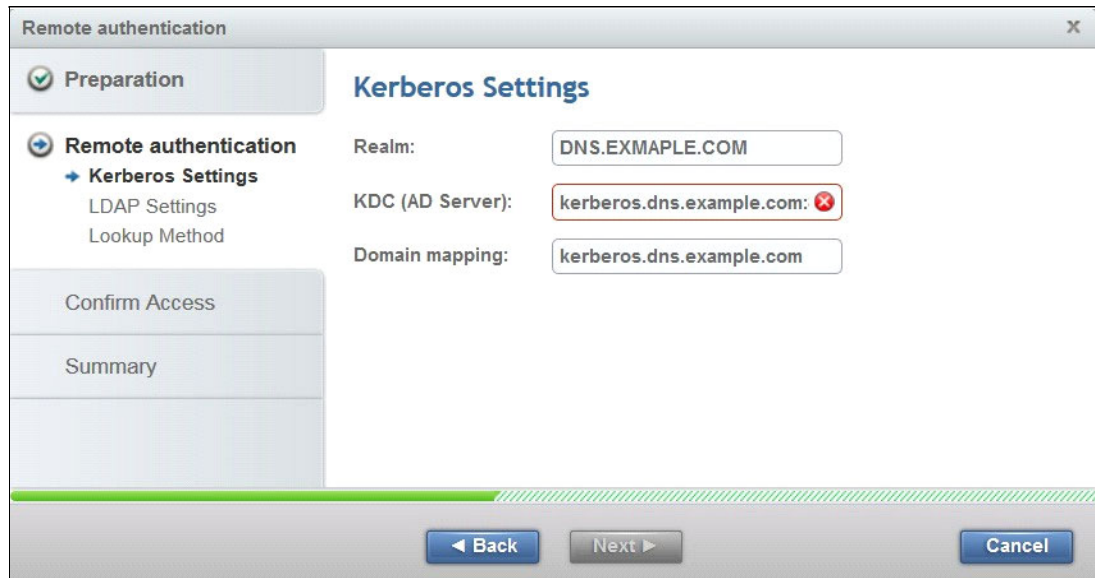


Figure 5-124 Kerberos settings example

LDAP

If LDAP is selected, the LDAP Settings window opens, as shown in Figure 5-125. Complete the required fields in this frame and press **Next**.



Figure 5-125 LDAP settings window

Configure the following settings on the Diamondback tape library:

- ▶ LDAP Repository URI. This field starts with `ldap://` or `ldaps://` and ends with a port number. The 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 Transport Layer Security (TLS) layer. Use `ldap://`. 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 look up the correct service credential user. As shown in Figure 5-126 on page 141, the user credentials are:

```
CN=Tom,CN=Users,DC=ldapsver1,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 Diamondback user does not have sufficient privileges to query the LDAP server, the Diamondback 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 checkmark 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.

The screenshot shows a 'Remote authentication' window with a sidebar on the left containing 'Preparation', 'Remote authentication', 'LDAP Settings', 'Lookup Method', 'Confirm Access', and 'Summary'. The main area is titled 'LDAP Settings' and contains the following fields:

- LDAP repository URI: (with a green checkmark)
- Secondary LDAP repository URI:
- LDAP StartTLS: Enable
- LDAP TLS certificate: (with a folder icon)
- Service Credentials (Optional):
 - User Name:
 - Password:

At the bottom, there are three buttons: 'Back', 'Next', and 'Cancel'.

Figure 5-126 LDAP settings example

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 DN's 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 5-127 on page 142 shows an example of a completed Simple lookup method page that includes an example of the format of the required fields.

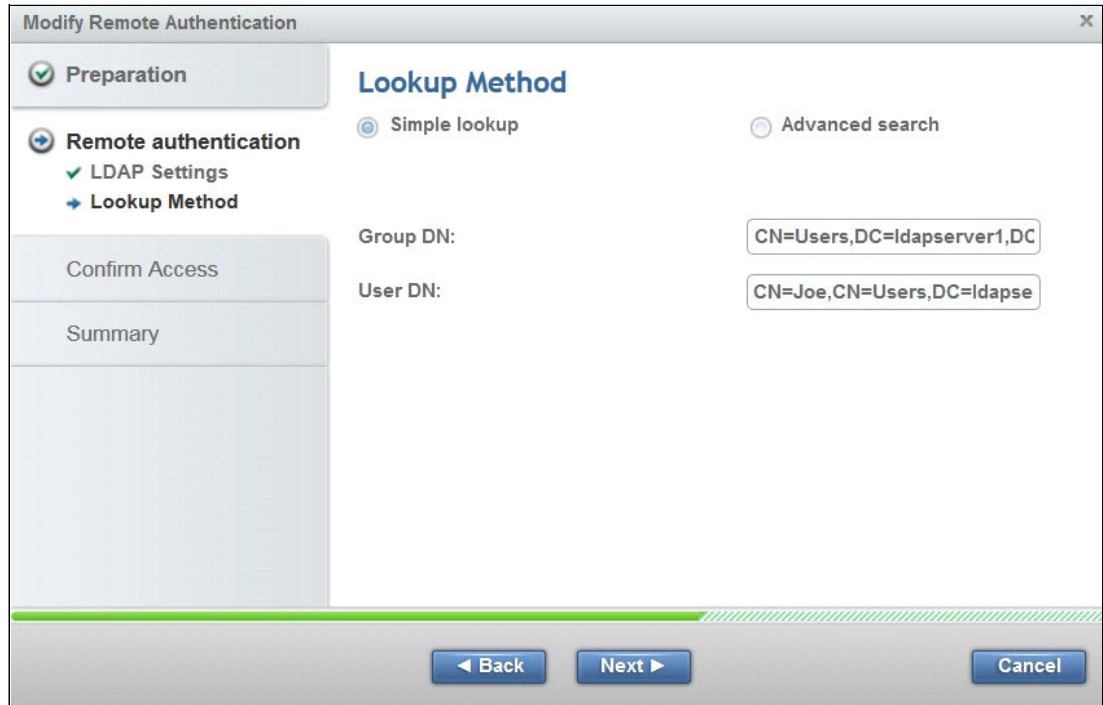


Figure 5-127 Simple lookup method window

Advanced search

The Advanced search 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 Diamondback 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. The 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 5-128 on page 143, which also shows how to complete the Advanced search within the LDAP lookup method.

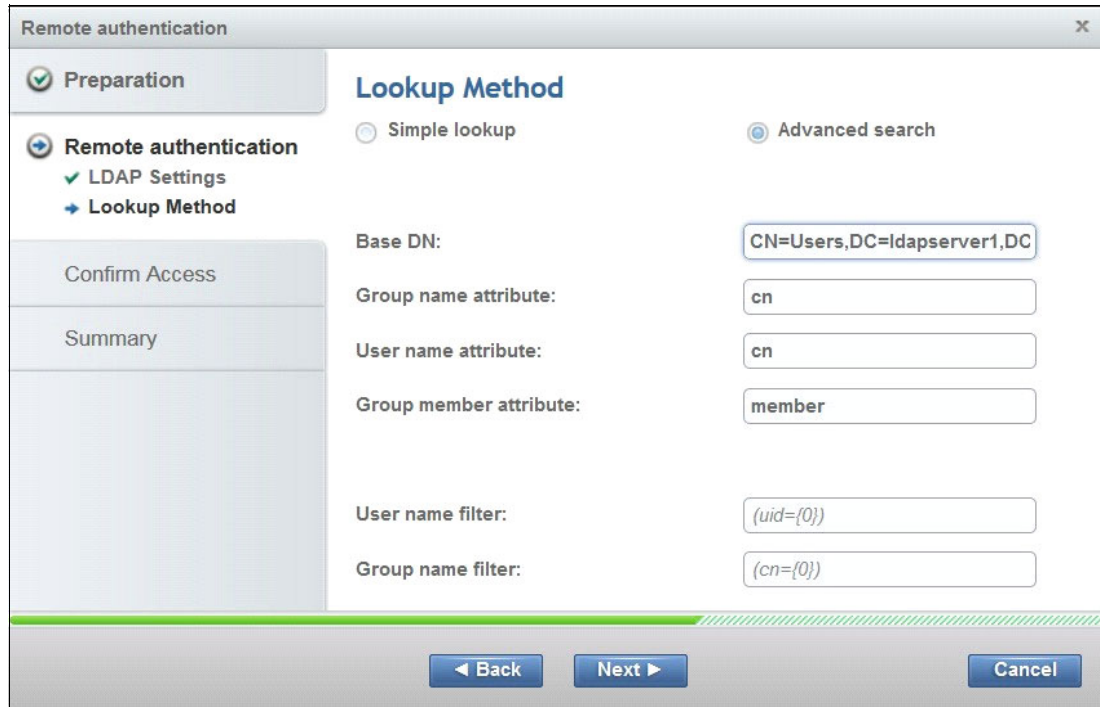


Figure 5-128 Advanced search settings window

After you enter the settings, click **Next** to continue to test that 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 an 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 is after 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 username must contain the full `racfid`: for example, `racfid=USER1,profiletype=USER,cn=RACF`. For more information, see the proceeding shaded box.
- ▶ **User name and Group name filters.**

Note: User name and Group name 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 login process.

When the user and 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. The following examples are 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, click **Security Remote Authentication Actions** to prompt 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 5-129 shows the Password and Session Policy window, which displays the current settings.

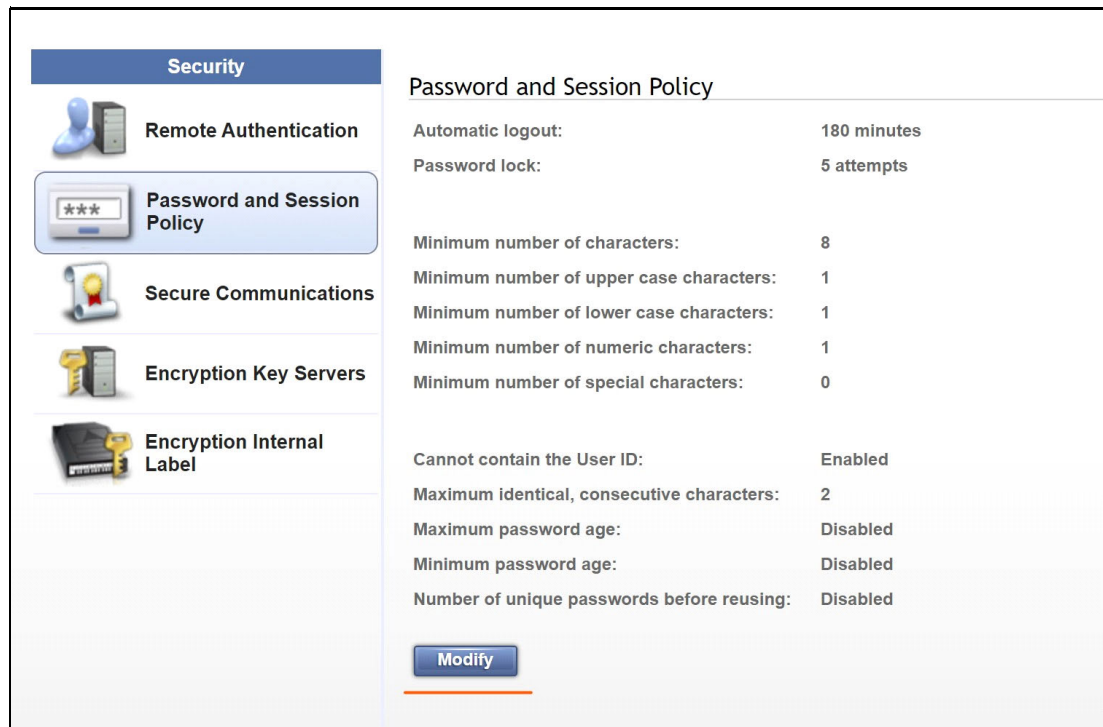


Figure 5-129 Password and session policy panel

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 5-130. These settings can be set back to the default settings by selecting **Reset to Default** in this window.

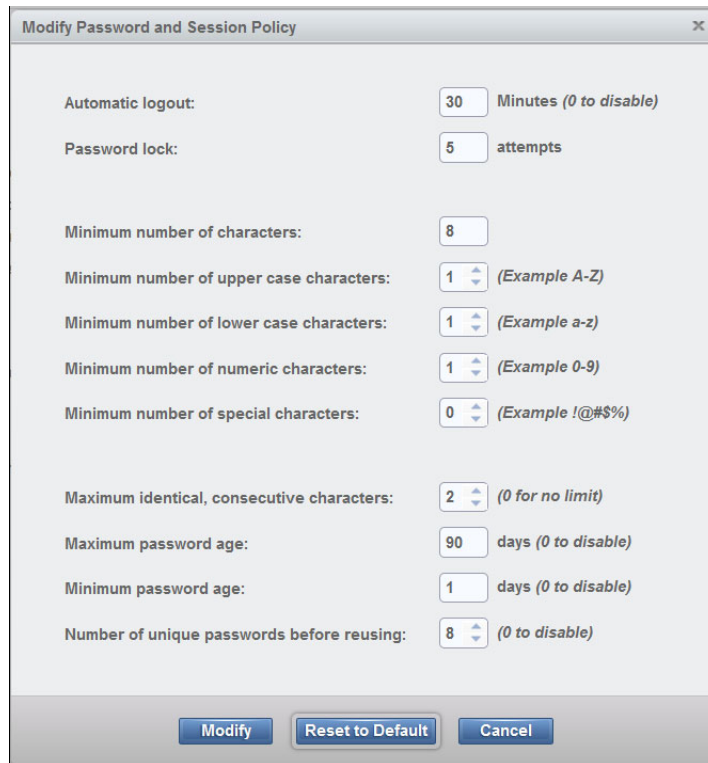


Figure 5-130

Secure communications

The IBM Diamondback tape library is secured with a Secure Sockets Layer (SSL). SSL is a protocol for encrypted (secure) transmission through the internet. The TLS 1.2 protocol is enabled by default. Meaning that any data that is exchanged between the Diamondback tape library and the browser is encrypted. In the browser, you must add and confirm a security exception to login. Use the Secure Communications page to configure the SSL settings (Figure 5-131 on page 146).

Selecting **Update Certificate** 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.

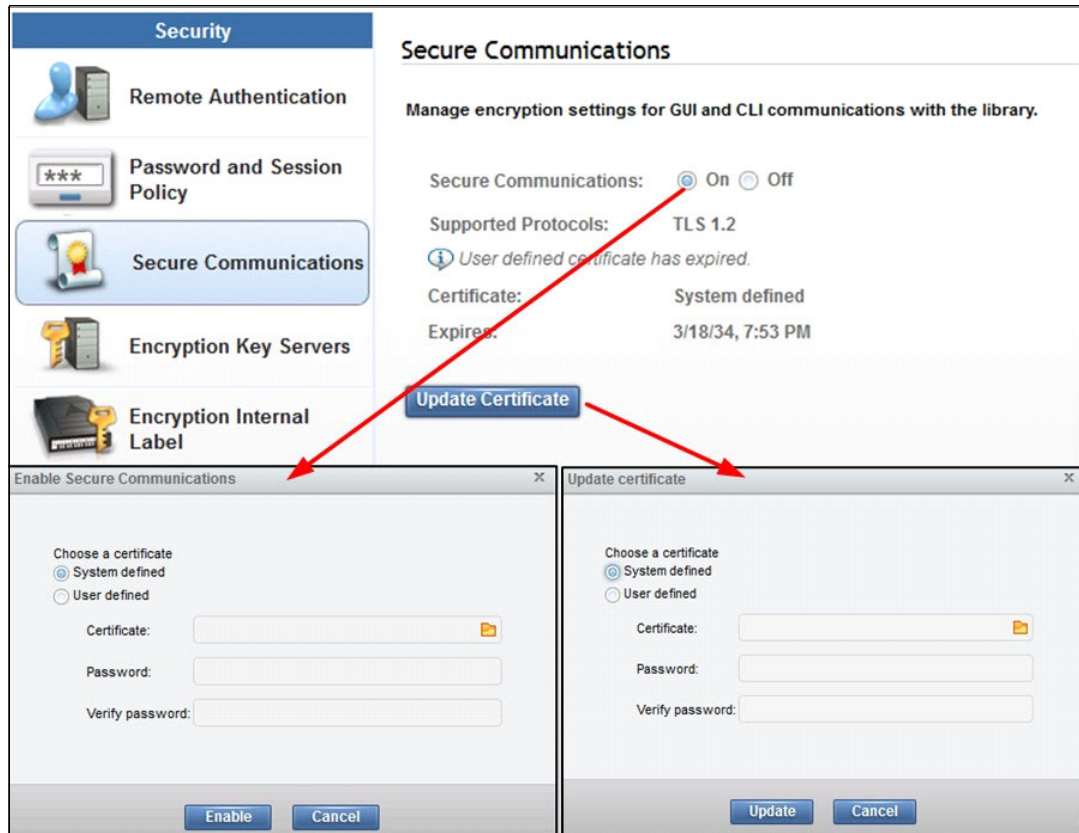


Figure 5-131 Secure Communications/HTTPS settings window

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. The certificate 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.

When you use CA-signed certificates on a Diamondback tape library that is configured with multiple IP addresses, it is recommended to include each IP address 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 [Getting started with SSL certificates](#).

Encryption key servers

Use the Encryption Key Servers page to manage which key servers are used to provide an encryption key to the library and drives.

Beginning with Firmware version 2.9.x, secure communication between the library and key server is supported. Certificates used are required to be in PEM-Format.

From the Add Encryption Key Server tab, define each encryption key server to be used with the Diamondback library as shown in Figure 5-132 on page 148. This action makes the IP address of the key server available when you set up library-managed encryption and can be assigned to each logical library as needed.

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 5.3.1, “Creating a logical library” on page 81 and Chapter 4, “Security and encryption” on page 49.

The Add Encryption Key Server window requires a check with an Internet Control Message Protocol (ICMP) ping to ensure that the server is available when added. Click **Ping** (Figure 5-132 on page 148) when you add a server to the list. Encryption key servers can also be added when you modify a logical library.

Note: The Diamondback supports the configuration of up to four encryption key servers for each logical library.

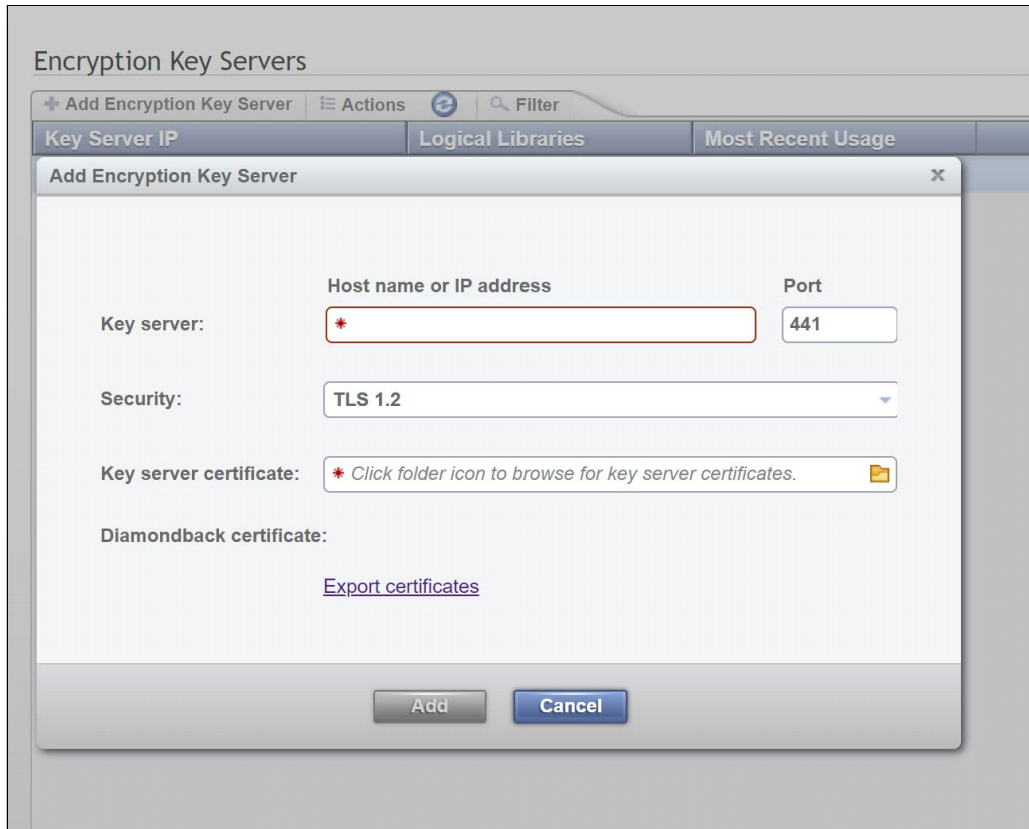


Figure 5-132 Add encryption key server window

Run Diagnostics option

To verify all of the network-attached devices are working, run diagnostics on the devices. Select **Actions** → **Run Diagnostics**. Then, select the drive or select drives to test and select **Run Diagnostics**.

The diagnostics process runs for a few minutes. When the diagnostics are completed, the window shows a check mark for devices that passed the test, as shown in Figure 5-133 on page 149, or a red error icon for devices that failed the test.

Note: When you use the Run Diagnostics function, a warning icon might appear in the Ethernet column if Internet Control Message Protocol (ICMP) requests are disabled.

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. The key server must be assigned to a logical library for diagnostics to run.

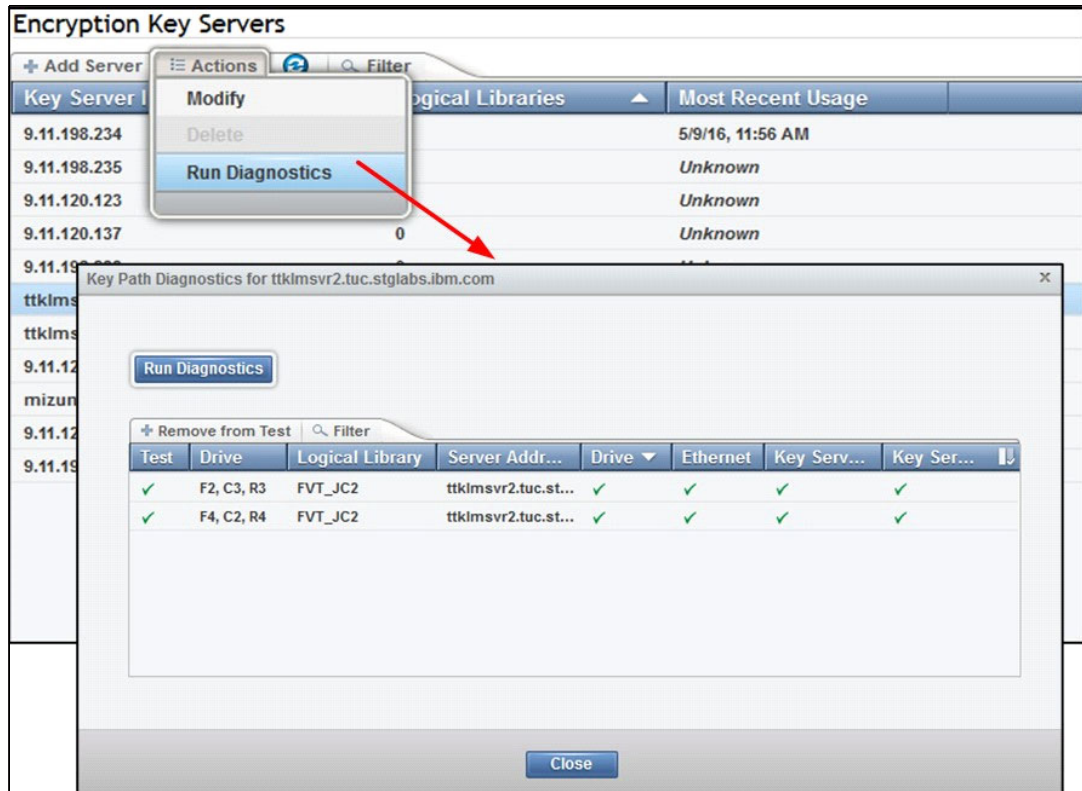


Figure 5-133 Encryption key servers diagnostics window

ICMP requests

Internet Control Message Protocol (ICMP) is a messaging protocol external to the Diamondback 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.

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 Diamondback tape library. Mapping the cartridge key labels to the key-encrypting labels enables the Diamondback 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.



Figure 5-134 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 5.3.1, “Creating a logical library” on page 81 and Chapter 4, “Security and encryption” on page 49.

Configuring Encryption Internal Label

To configure key label mapping, select **Encryption Internal Label** and then select the **Create Mapping** tab, as shown in Figure 5-135. Select the method to use and the key labels, as required.

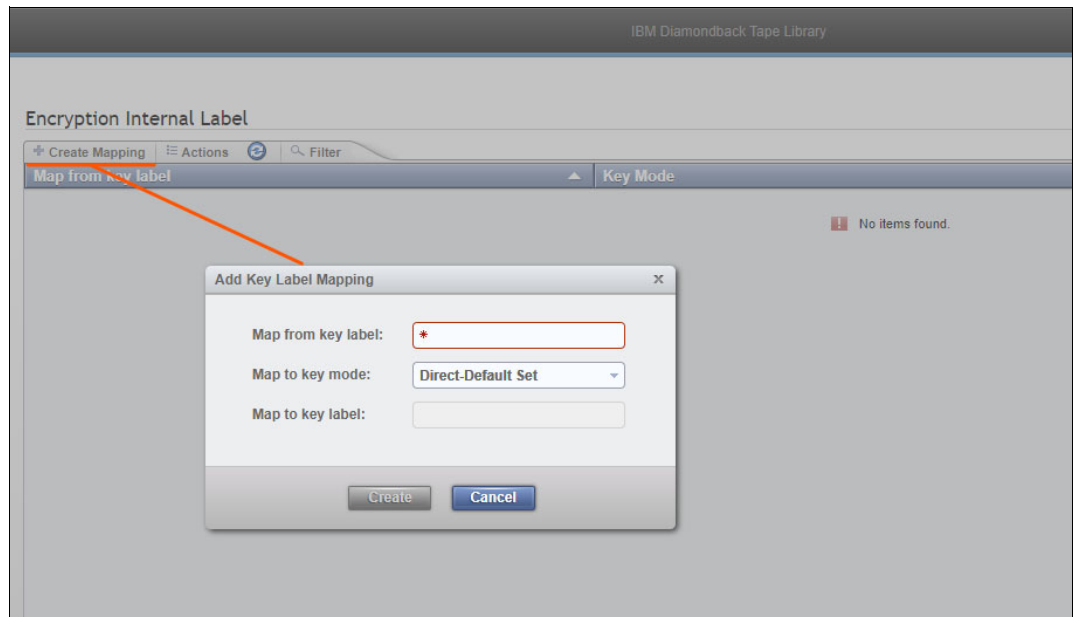


Figure 5-135 Create Encryption label panel

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.
- ▶ Direct-Specific. The specified key label references a symmetric data key.

5.7.4 GUI Preferences page

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 Diamondback Library to enable or disable the animation of the navigation dock in the GUI. Figure 5-136 on page 151 shows an example of the GUI with animation enabled.

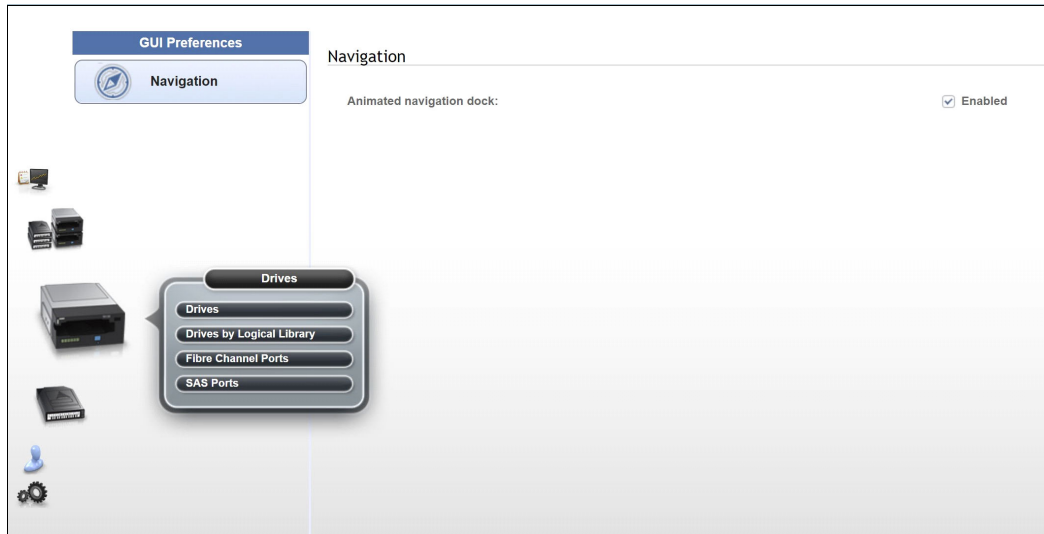


Figure 5-136 Navigation menu with animation enabled

Animations can be disabled if the enlarged icons obscure any information in the window. Turning off the navigation animation can help increase performance of the GUI when remotely connecting to the Diamondback over a slow network connection. Figure 5-137 shows example of the GUI with animation disabled.

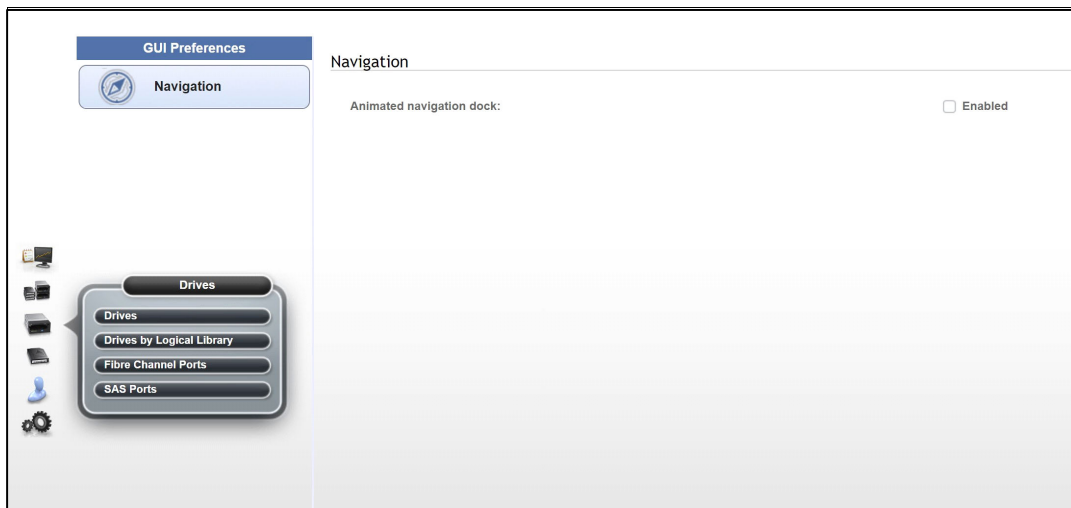


Figure 5-137 Navigation menu with animation disabled

5.8 Remote support

Remote support for the Diamondback tape library involves the use of a Call Home feature to detect, report, and solve problems. Because the Diamondback library does not have an IMC-Integrated Management Console (IMC) like the TS4500, an external TSSC is required.

Figure 5-138 on page 152 shows a diagram how the external TSSC is connected to the Diamondback library.

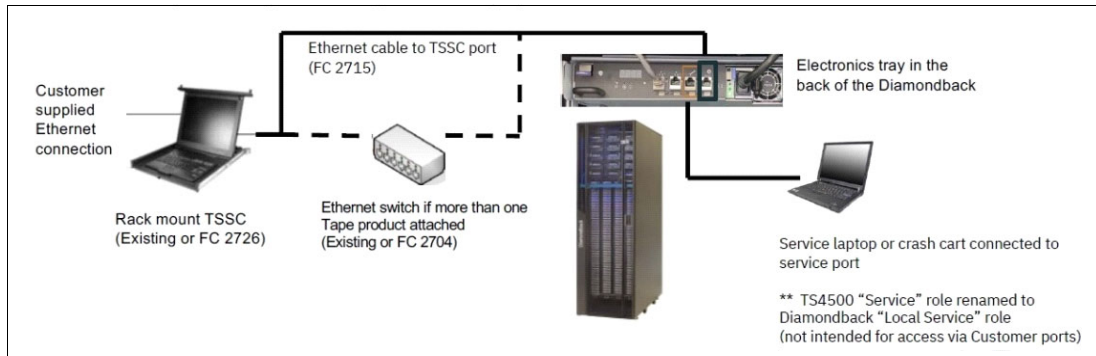


Figure 5-138 Diamondback TSSC configuration example

Remote support is available for the Diamondback tape library through its Call Home capability. The Call Home feature uses an external TSSC 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 physical location of the library. The library can also periodically send support information, such as configuration information, code versions of the library and drives, and error logs to IBM.

The Call Home feature of the Diamondback tape library has three different but related capabilities: Problem Call Home, Heartbeat Call Home, and Test Call Home. The Diamondback 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 drives
- ▶ Model and serial number of library
- ▶ Type and serial number of each drive
- ▶ Code version of the library and each drive
- ▶ Any machine-detectable features.

If a Problem Call Home is initiated, the library also sends the tape library logs and drive logs that relate to the problem, and it generates a Cognitive Support Platform (CSP) case number.

Call Home Database

CSP, which is used to manage the problem records (cases) between the various support levels and drive them to resolution, includes high-level error information, which is provided to the support team. They can then provide the most common solutions to the problem. A system expert might automatically update the CSP Case with the recommended service actions based on the cause of the call home. Additional detailed information is then placed into ECuRep and processed by the DRI Call Home DB which further aids problem determination if needed or can be mined to analyze trends and configuration information.

Problem Call Home

When the Diamondback tape library or one of its drives detects a problem, the library initiates a Call Home operation. This Call Home operation creates a CSP case with details on the issue reported. These pages 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) or customer-replaceable-units (CRUs), that might be needed to address the issue. If a client

suspects that there is a problem with the Diamondback library, a transfer of logs to IBM can be initiated from the GUI in parallel to an IBM support ticket being opened either by phone (1800-IBM-SERV) or through the [IBM support portal](#).



Figure 5-139 Log export actions

Heartbeat Call Home

The Diamondback tape library sends the Heartbeat Call Home daily to ensure that the Call Home function operates correctly. The Heartbeat Call Home is sent once per day, 1 hour after a power cycle, and 1 hour after a code update completes. Figure 5-140 shows how to enable Call Home.

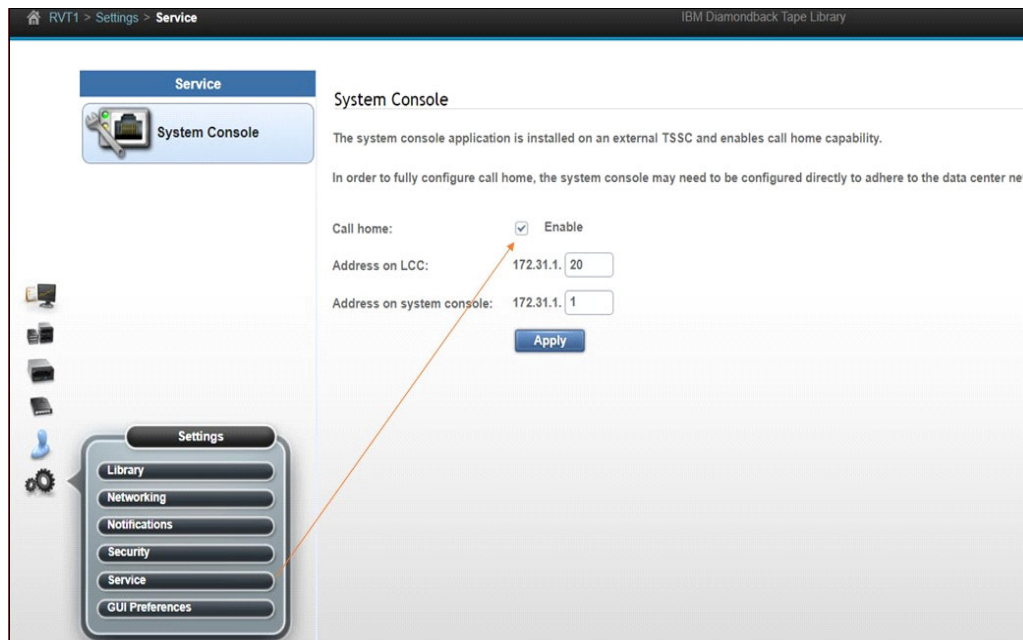


Figure 5-140 Library Call home configuration

Test Call Home

When the SSR services the library, the SSR can issue a Test Call Home operation to IBM from the Diamondback management GUI. Using the Test Call Home function is a good way to test the environment to confirm if the required Ethernet ports are open in the network.

5.8.1 Remote support by using an external TSSC

The only method that is used to perform an Electronic Customer Care (ECC) Call Home operation is through the external TSSC over a broadband connection. The Diamondback tape library sends Call Home information across an Ethernet connection to the external TSSC.

The TSSC then performs the ECC Call Home operation and sends the information to IBM through the Ethernet (broadband) connection provided by the customer. Table 5-6 lists the capabilities of remote support when an external TSSC is attached to the library.

Table 5-6 Support and call home function availability

Location	Event Support	Support
Client site	Call home event	<ul style="list-style-type: none"> ▶ Error initiated ▶ Heartbeat (regular interval) ▶ Test
	System console support capability	<ul style="list-style-type: none"> ▶ 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 the Fix Central Fix Distribution Library)
	System console and remote support service tools	<ul style="list-style-type: none"> ▶ Code image broadcast ▶ Call Home event log review ▶ End-of-call completion report
IBM Support	System console remote access	<ul style="list-style-type: none"> ▶ Simultaneous call in and Call Home Data transmission (Transmission Control) ▶ Protocol/Internet Protocol (TCP/IP) supported
	IBM Call Home Database	<ul style="list-style-type: none"> ▶ Always available (24x7) access by IBM Support staff ▶ Error analysis and search capability

5.8.2 Remote support security

Security considerations with the Diamondback management GUI

Remote support security for the Diamondback tape library is performed through a system console when using the management GUI is described.

The management GUI does not allow access to client data, and it does not allow 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 loading firmware that is not the latest release level.

These security concerns can be addressed by using the password, user role, and authentication features that are provided by the management GUI.

Library management by using an external TSSC

The external TSSC also provides optional Ethernet-based 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 5-7. You can modify the port numbers by using the Diamondback management GUI.

Table 5-7 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

Port	Type of data	Direction	Protocol
636	LDAPS	Bidirectional	TCP and UDP
443	HTTPS	Bidirectional	TCP

Data Privacy

IBM intends to protect your personal information and to maintain its integrity. IBM implements reasonable physical, administrative, and technical safeguards to help protect your personal information from unauthorized access, use, and disclosure. For example, the Diamondback tape library provides IBM data about the asset usage and configuration only and does not reflect private use of the asset. When diagnostics are sent to IBM, in the case of a problem submission, that data is routed directly to a secured infrastructure. Only individuals with a need to know are given access while working to resolve your problem.

Dial-out security features of the external TSSC

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 disabled. The Call Home feature can be enabled by using the web management interface.

None of the client data that is stored on the tape or in memory for the Diamondback tape library is transmitted or accessed in a Call Home session. Call Home is enabled or disabled through the service menu of the Diamondback management GUI.

Dial-in security features of the external TSSC (Assist On-site/AOS)

Dial-in is used by IBM support to log on to the TSSC console and provide service support. This feature must be enabled by the client or SSR. 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.

Security and privacy are fundamental concerns when granting remote access to corporate IT assets. Assist On-site uses the latest security technology to ensure that the data exchanged between IBM Support engineers and clients is secure. Identities are verified and protected with industry-standard authentication technology, and Assist On-site sessions are kept secure and private with the use of randomly generated session keys and advanced encryption.

Assist On-site allows IBM Support engineers to remotely access client's computers to identify and resolve technical issues in real time. Assist On-site facilitates problem determination and remediation by providing a powerful suite of tools that enables IBM Support engineers to quickly complete root cause analysis and take appropriate corrective action.

When a screen-sharing session begins, the Support engineer is connected to the client's computer (here TSSC/IMC) via a relay server. Long, randomly generated session keys are issued to both participants to ensure that only the designated parties are connected. During the session, all transferred information, including screen views, file-transfer data and identities, are encrypted. Encryption and decryption are from end to end, so data can't be intercepted during transit and can be viewed only by using the Assist On-site console.

Assist On-site sessions are protected by strong password authentication. Support engineers are authenticated by using a challenge and response password exchange.

Assist On-site implements outbound connections, which are protected by TLS encryption over a port 443 session to prevent intruder access to the information exchanged during all Assist On-site sessions.

Chat, screen viewing, screen-sharing, and file transfer data are encrypted end to end and packets are never decrypted in transit by the communication servers.

Assist On-site works seamlessly with most firewalls. Usually, Assist On-site connections are possible without any firewall reconfiguration. Assist On-site requires access to outbound ports at both ends of a connection, so there is no need to open holes in firewalls.

Assist On-site automatically chooses the relay server that provides the best end-to-end performance. All relay servers are available from all geographies, with performance typically better from the relay server closest to the client system.

Note: The Diamondback library does not support dial-in directly, a connection must be made through an attached TSSC.

The external TSSC supports the following data security features 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 Diamondback 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.

Port information for firewall environments

Table 5-8 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 5-8 Diamondback 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	441	Outbound	SSL

Function	Port	Direction (from library)	Protocol
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 external TSSC 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.

5.8.3 Saving and restoring the Diamondback Backup to the TSSC

If there is an external TSSC attached to the Diamondback library and the TSSC firmware level is at 9.4.10 or later, the TSSC can be used for making a backup of the Diamondback library. In cases where more than one Diamondback library is attached, a backup is made of all attached Diamondback libraries.

IBM recommends taking a backup of any Diamondback library configuration regularly either manually or by using the automatic backup function of the TSSC. The TSSC offers the following options for an automated backup:

- ▶ Not scheduled
- ▶ Daily
- ▶ Weekly
- ▶ Monthly

Making a backup of the configuration

The following steps show how to configure a backup and perform a restore of the Diamondback by using the TSSC:

1. Log in to TSSC, see Figure 5-141.

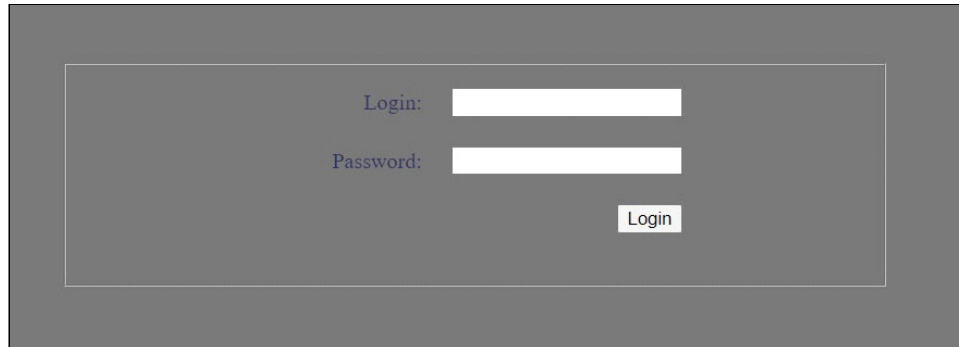


Figure 5-141 TSSC login screen

2. Click **Backup/Restore** (Figure 5-142).



Figure 5-142 Backup/Restore TSSC option

3. On the next window, click **Backup/Restore Library Configuration Data**. (Figure 5-143).

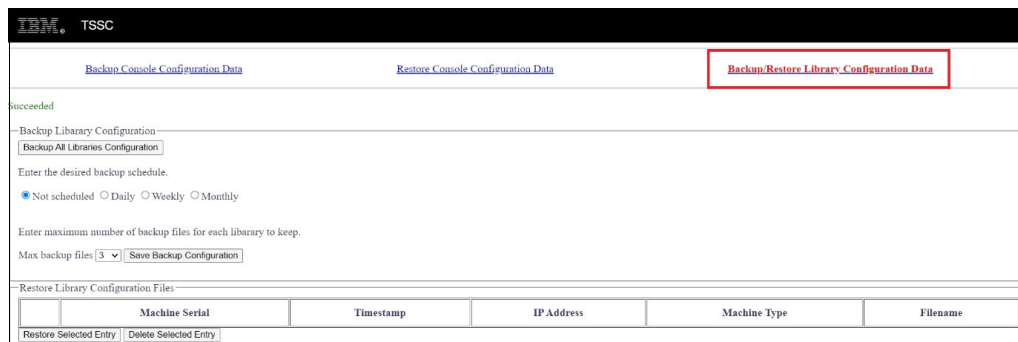


Figure 5-143 Backup/Restore configuration data

4. The next step allows the user to create a backup schedule or to back up all libraries. Choose either not scheduled, Daily, Weekly, or Monthly options to create a backup schedule. To manually run a backup of all attached libraries, use the **Backup All Libraries**

button. When clicked, a backup is started for all attached Diamondback libraries, and the table is updated.(Figure 5-144 and Figure 5-145).

Succeeded

Backup Library Configuration

Backup All Libraries Configuration

Enter the desired backup schedule.

Not scheduled Daily Weekly Monthly

Enter maximum number of backup files for each library to keep.

Max backup files 3

Restore Library Configuration Files

	Machine Serial	Timestamp	IP Address	Machine Type	Filename
<input type="radio"/>	13	2023/06/09 11:12:17	172.31.1.20	DIAMONDBACK	DIAMONDBACK_B0023_2.9.0.3-10L.00_SAVECONF_20230609111217.dbz

Figure 5-144 Backup library configuration options and backup configuration button

Allow a couple of minutes for the backup to finish. The backup files are shown in the Restore Library Configuration Files section.

	Machine Serial	Timestamp	IP Address	Machine Type	Filename
<input type="radio"/>	13	2023/06/09 11:12:17	172.31.1.20	DIAMONDBACK	DIAMONDBACK_B0023_2.9.0.3-10L.00_SAVECONF_20230609111217.dbz

Figure 5-145 Configuration backup version details and file names

When one of the scheduled options of Daily, Weekly, or Monthly is used the change must be confirmed by clicking **Save Backup Configuration** (Figure 5-146).

Backup Library Configuration

Backup All Libraries Configuration

Enter the desired backup schedule.

Not scheduled Daily Weekly Monthly

Enter maximum number of backup files for each library to keep.

Max backup files 3

Figure 5-146 Save backup configuration selection

The save of the backup configuration is confirmed with a **Succeeded** message in the web GUI. See Figure 5-147 on page 161.

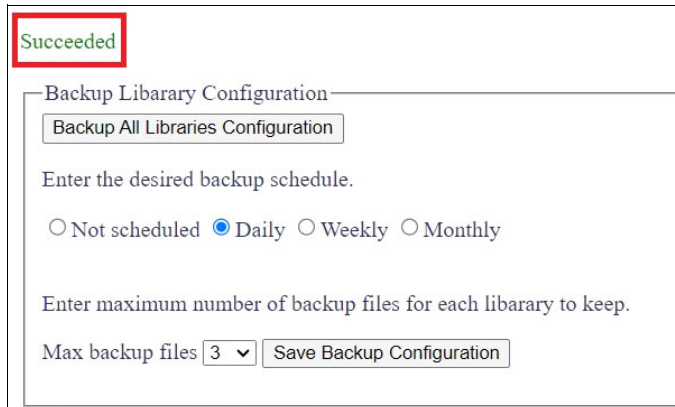


Figure 5-147 Successful library configuration backup

Restoring the configuration

For the restore option, one of the backup files must be selected. The backup file is uploaded to the library and takes approximately 10 minutes to complete.

After the upload of the backup file to the library, the library is rebooted.

Note: Before you start the restore action, stop all I/O from the host because when the restore is complete the library is rebooted.

Use the following steps to restore a library backup configuration file:

1. Select the file that you want to restore (Figure 5-148).
2. Click the **Restore Selected Entry**.

Restore Library Configuration Files					
	Machine Serial	Timestamp	IP Address	Machine Type	Filename
<input checked="" type="radio"/>	13	2023/06/09 11:12:17	172.31.1.20	DIAMONDBACK	DIAMONDBACK_B0023_2.9.0.3-10L.00_SAVECONF_2023060911121

Below the table, there are two buttons: 'Restore Selected Entry' (highlighted with a red box) and 'Delete Selected Entry'.

Figure 5-148 Select and restore a library configuration file

After a successful restore, a **Succeeded** message is displayed in the web GUI (Figure 5-149).




Figure 5-149 Restore succeeded message

5.9 IBM Net Promoter Score and Feedback

Net Promoter Score (IBM NPS®) is a tool that is used to gauge a firm's customer relationship with IBM. It serves as an alternative to traditional customer satisfaction research. NPS is widely used, and more than two-thirds of Fortune 1000 companies use the metric.

Every user of the Diamondback GUI is contacted after 30 days of use to provide feedback about the Diamondback by way of a *Suggested Task*. After 30 days, the Suggested Task is shown (Figure 5-150 and Figure 5-151).




Figure 5-150 Status icon showing pending task

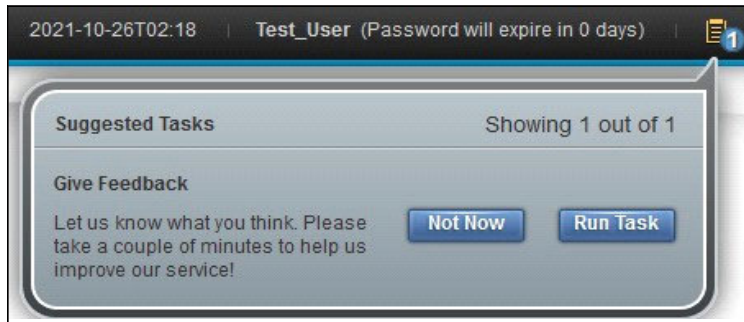


Figure 5-151 Pending feedback task

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** (Figure 5-152).

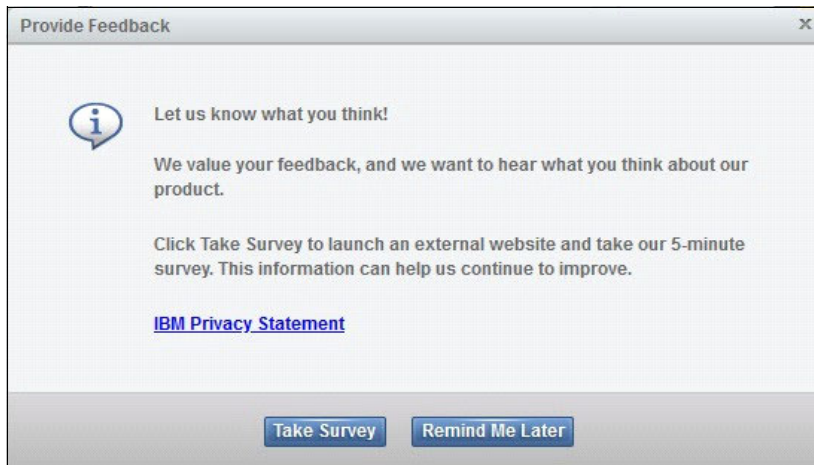


Figure 5-152 Provide feedback task window

Clicking **Take Survey** opens the Medallia web page where you can provide feedback.

Note: You can also review the IBM Privacy Statement by clicking the included link.



IBM Diamondback REST API

This chapter describes the Diamondback REST API.

The Diamondback REST API can be used to obtain information from and send commands to the library.

At the time of writing, there are two methods for sending REST requests to the library. The first method uses in-band SCSI commands through a Control Path drive that is issued from an attached host. This is called REST over SCSI, or RoS. The second method is the more traditional REST API approach of using an Ethernet network connection to communicate with a webserver on the library and issue requests through a URL. This is referred to as REST over Ethernet, or RoE.

RoS allows library administrators to send REST requests by issuing a SCSI Write Buffer command with a URL and receive an HTTP response with a Read Buffer command. If RoS is enabled, which is the default setting, the library accepts and responds to commands. If it is disabled, RoS requests are rejected. These commands can be sent to the library by using the `sg3_utils` package for Linux or ITDT for multiple operating systems.

RoE functions similarly to any other REST API in that it is run over an Ethernet network, and standard tools for interacting with a REST API work with it. The only major difference between RoE and RoS commands is that RoE requires a user to be authenticated before any commands are accepted by the library.

Some of the attributes in this API might seem redundant or unnecessary. This is because it is intended to be compatible with any software that uses the TS4500 REST API.

This is meant as a general guide to Diamondback REST API usage. For the most up-to-date information, see [Diamondback Tape Library REST over SCSI API](#).

This chapter includes the following topics:

- ▶ 6.1, “REST API Overview” on page 58
- ▶ 6.2, “Using `sg3_utils` (Linux)” on page 105
- ▶ 6.3, “Using ITDT for RoS” on page 106
- ▶ 6.4, “REST over Ethernet” on page 108
- ▶ 6.5, “HTTP return codes” on page 110

6.1 REST API Overview

6.1.1 Commands

REST API commands are issued along with a URL endpoint and any parameters that are needed. REST API commands fit into the following categories:

- ▶ GET commands that request information from the library.
- ▶ POST commands that initiate a task (such as inventorying the library) or an action (resetting a drive).
- ▶ PUT commands that change a library setting (such as changing the library's name). These commands were previously listed as PATCH commands in TS4500, but this has been changed with the introduction of REST over Ethernet.

6.1.2 Endpoints

Diamondback REST API endpoints use JSON structures to return data to the requester. Due to the nature of JSON parsers, the attributes in the following examples might not be in the exact order that is specified by their associated tables because they are formatted to be easier to read.

Accessor

The following endpoints are available for accessors:

- ▶ **GET /v1/accessors**

Retrieve a list of accessors with status and attributes for monitoring purposes. Note that only one accessor is ever listed for a Diamondback library.

Example 6-1 Get /v1/accessors output

```
<Response [200]>
[
  {
    "location": "accessor_Aa",
    "state": "onlineActive",
    "stateReferenceEvent": null,
    "driveAccess": "normal",
    "cartridgeAccess": "normal",
    "pivots": 0,
    "barCodeScans": 42214,
    "velocityScalingXY": 100,
    "velocityScalingPivot": 100,
    "travelX": 654,
    "travelY": 82985,
    "getsGripper1": 4946,
    "putsGripper1": 4946,
    "getsGripper2": 4372,
    "putsGripper2": 4368,
    "temperature": 22.4,
    "humidity": 63.7
  }
]
```

Table 6-1 Get accessors attribute details

Attribute	Details
location (string)	Defines which accessor this is. Always set to accessor_Aa.
state (string)	The current state of the accessor. Possible values include inServiceMode, noMovementAllowed, bothGrippersFailed, gripper1Failed, gripper2Failed, scannerFailed, noMotorPower, calibrating, onlineActive. The default state of a working accessor is onlineActive.
stateReferenceEvent (number)	The event ID of the event that caused a state change. If no error or warning event caused this state, then the value is <i>null</i> .
driveAccess (string)	Indicates whether the accessor can reach a drive or not.
cartridgeAccess (string)	Indicates whether the accessor can reach all cartridges or not.
pivots (number)	Total number of pivots the accessor has performed. This value will always be 0 in a Diamondback library because the accessor does not pivot.
barcodeScans (number)	Total number of barcodes the accessor has scanned.
velocityScalingXY (number)	A percentage (in increments of 10) that represents the velocity of the accessor.
velocityScalingPivot (number)	A percentage (in increments of 10) that represents the pivot velocity. This is not used by Diamondback libraries.
travelX (number)	Number of meters the accessor has traveled horizontally in its lifetime.
travelY (number)	Number of meters the accessor has traveled vertically in its lifetime.
getsGripper1 (number)	Number of times the first gripper has been engaged to retrieve a cartridge in its lifetime.
putsGripper1 (number)	Number of times the first gripper has been engaged to place a cartridge in a slot or drive in its lifetime.
getsGripper2 (number)	Number of times the second gripper has been engaged to retrieve a cartridge in its lifetime.
putsGripper2 (number)	Number of times the second gripper has been engaged to place a cartridge in a slot or drive in its lifetime.

Attribute	Details
temperature (number)	The most recent temperature measurement from the accessor in Celsius to one decimal point. This will return <i>null</i> if it was unable to be measured.
humidity (number)	The most recent humidity percentage measurement from the accessor to one decimal point. This will return <i>null</i> if it was unable to be measured.

► **GET /v1/accessors/<location>**

Get status and attributes for an accessor in a specific location. This is the same output as the *GET /v1/accessors* command on a Diamondback library.

Example 6-2 GET /v1/accessors/<location> out

```
<Response [200]>
[
  {
    "location": "accessor_Aa",
    "state": "onlineActive",
    "stateReferenceEvent": null,
    "driveAccess": "normal",
    "cartridgeAccess": "normal",
    "pivots": 0,
    "barCodeScans": 42214,
    "velocityScalingXY": 100,
    "velocityScalingPivot": 100,
    "travelX": 654,
    "travelY": 82985,
    "getsGripper1": 4946,
    "putsGripper1": 4946,
    "getsGripper2": 4372,
    "putsGripper2": 4368,
    "temperature": 22.7,
    "humidity": 63.6
  }
]
```

► **PUT /v1/accessors/<location>/velocityScaling {"velocityScalingXY": <percentage>}**

Sets the velocity scaling for the robotic accessor's vertical and horizontal movement. This must be sent in increments of 10, or it will be rejected by the library.

A successful update with this command returns an RC 201 message.

Authentication

The following commands relate to authentication and user accounts:

► GET /v1/authentication/passwordPolicy

Retrieves details about the password policy for the library's web GUI and REST over Ethernet.

Example 6-3 GET /v1/authentication/passwordPolicy output

```
[
  {
    "passwordAttemptsAllowed": 5,
    "minChars": 8,
    "minUpperCaseChars": 1,
    "minLowerCaseChars": 1,
    "minNumericChars": 1,
    "minSpecialChars": 1,
    "maxIdenticalConsecutiveChars": 2,
    "maxPasswordAge": 90,
    "minPasswordAge": 1,
    "uniquePasswords": 8
  }
]
```

Table 6-2 GET authentication/passwordPolicy details

Attribute	Details
passwordAttemptsAllowed (number)	Number of failed logins before a user account is locked.
minChars (number)	Minimum number of characters required for a new password.
minUpperCaseChars (number)	Minimum number of upper-case characters required in a new password.
minLowerCaseChars (number)	Minimum number of lower-case characters required in a new password.
minNumericChars (number)	Minimum number of numeric characters required in a new password.
minSpecialChars (number)	Minimum number of special characters (! @ # \$ %) required in a new password.
maxIdenticalConsecutiveChars (number)	Maximum number of repeating characters allowed in a new password.
maxPasswordAge (number)	Number of days before a password will expire and require a user to reset it.
minPasswordAge (number)	Number of days that must pass before a password is allowed to be changed.
uniquePasswords (number)	Number of unique passwords required before a password can be reused.

► **GET /v1/authentication/roles**

Returns a list of user roles and their details from the library.

Example 6-4 GET /v1/authentication/roles output

```
[
  {
    "name": "Administrator",
    "type": "default",
    "users": 2,
    "sessions": 0
  },
  {
    "name": "Superuser",
    "type": "default",
    "users": 0,
    "sessions": 0
  },
  {
    "name": "Monitor",
    "type": "default",
    "users": 2,
    "sessions": 0
  },
  {
    "name": "Local Service",
    "type": "default",
    "users": 3,
    "sessions": 0
  }
]
```

Table 6-3 Get authentication/roles details

Attribute	Details
name (string)	The name assigned to the role.
type (string)	Denotes a system-defined user role as <i>default</i> or a user-defined role as <i>custom</i> .
users (number)	Number of users with this role.
sessions (number)	Number of active sessions using this role.

► **GET /v1/authentication/roles/<name>**

Returns the same attributes as *GET /v1/authentication/roles* for a single role.

Example 6-5 GET /v1/authentication/roles/<name> output

```
[
  {
    "name": "Administrator",
    "type": "default",
    "users": 2,
    "sessions": 0
  }
]
```

```
}  
]
```

► **GET /v1/authentication/sessions**

Returns a list of active user sessions on the library and their details.

Example 6-6 GET /v1/authentication/sessions output

```
[  
  {  
    "name": "admin",  
    "authenticationMethod": "local",  
    "role": "Administrator",  
    "address": "9.61.250.239",  
    "lastConnected": "2023-08-21T11:38:13-0700"  
  }  
]
```

Table 6-4 GET /v1/authentication/sessions details

Attribute	Details
name (string)	Username for an active session.
authenticationMethod (string)	The authentication method used. Potential values are <i>local</i> , <i>ldap</i> , or <i>Kerberos</i> .
role (string)	The role associated with this username.
address (string)	IP address from which the session was initiated.
lastConnected (string)	The last time this user logged in.

► **GET /v1/authentication/sessions/<name>**

Returns the same attributes as *GET /v1/authentication/sessions* for a single user.

Example 6-7

```
[  
  {  
    "name": "admin",  
    "authenticationMethod": "local",  
    "role": "Administrator",  
    "address": "9.61.250.239",  
    "lastConnected": "2023-08-21T11:38:13-0700"  
  }  
]
```

► **GET /v1/authentication/userAccounts**

Returns a list of user accounts and their details from the library.

Example 6-8 GET /v1/authentication/userAccounts output

```
[  
  {  
    "name": "Service",
```

```

    "state": "disconnected",
    "role": "Local Service",
    "email": "",
    "lastConnected": "2023-08-21T11:20:53-0700"
  },
  {
    "name": "localGUI",
    "state": "disconnected",
    "role": "Monitor",
    "email": "",
    "lastConnected": "2023-01-25T15:25:05-0700"
  },
  {
    "name": "admin",
    "state": "connected",
    "role": "Administrator",
    "email": "",
    "lastConnected": "2023-08-21T11:38:12-0700"
  },
  {
    "name": "devService",
    "state": "disconnected",
    "role": "Local Service",
    "email": "",
    "lastConnected": "2023-08-16T17:07:43-0700"
  },
  {
    "name": "AARON",
    "state": "disconnected",
    "role": "Administrator",
    "email": "",
    "lastConnected": "2023-04-18T12:54:25-0700"
  },
  {
    "name": "andrea8A",
    "state": "disconnected",
    "role": "Monitor",
    "email": "andreadochoa@ibm.com",
    "lastConnected": null
  },
  {
    "name": "aaronService",
    "state": "disconnected",
    "role": "Local Service",
    "email": "",
    "lastConnected": null
  }
]

```

Table 6-5 GET /v1/authentication/userAccounts details

Attributes	Details
name (string)	Username for the account.

Attributes	Details
state (string)	Either <i>connected</i> or <i>disconnected</i> depending on whether the user is logged int currently.
role (string)	User role associated with this user.
email (string)	Email address associated with this user. This can be an empty string if there is no email address for the user.
lastConnected (string)	Timestamp of the last time the user logged in.

- ▶ **GET /v1/authentication/userAccounts/<name>**
Returns the same attributes as *GET /v1/authentication/userAccounts* for a single user.
- ▶ **POST /v1/authentication/userAccounts/<username>/setPassword {"password": <password>, "expirePassword": <"yes" | "no">}**
Sets the password for a specified user account.

```
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 0
```

Figure 6-1 Successful command update example

- ▶ **POST /v1/authentication/userAccounts {"role": <role>, "name": <username>, "password": <password>, "expirePassword": <"yes" | "no">, "emailAddress": <email>}**
Creates a new user account on the library with the specified attributes and sets the password. The email address attribute is optional. The user must change this password on first login.

```
HTTP/1.1 201 Created
Content-Type: application/json
Content-Length: 0
```

Figure 6-2 Successful command update example

Cleaning cartridges

The following commands are for the use of cleaning cartridges:

- ▶ **GET /v1/cleaningCartridges**
Returns a list of cleaning cartridges and their details from the library.

Example 6-9 GET /v1/cleaningCartridges output

```
[
  {
    "volser": "CLN041L1",
    "state": "importing",
    "accessible": "normal",
```

```

    "cleansRemaining": 50,
    "location": "slot_F1C3R5T1",
    "mediaType": "LTO",
    "mostRecentUsage": "2023-06-13T13:27:03-0700",
    "internalAddress": "FF030D"
  }
]

```

Table 6-6 GET /v1/cleaningCartridges details

Attribute	Details
volser (string)	Volume serial number or bar code for the cartridge.
state (string)	Health status of the cartridge. Potential values are <i>normal</i> , <i>importing</i> , <i>exportQueued</i> , <i>uncertainBarcode</i> , and <i>unknown</i> .
accessible (string)	Indicates whether the accessor can reach the cartridge.
cleansRemaining (number)	Number of cleans remaining before the cartridge must be replaced.
location (string)	Current location of the cartridge.
mediaType (string)	Indicates whether this is an LTO or 3592 cartridge.
mostRecentUsage (string)	The last time this cartridge was mounted to a drive or <i>null</i> if it has never been mounted.
internalAddress (string)	A unique identifier for the cartridge, used to differentiate cartridges in the event of duplicate volsers.

► **GET /v1/cleaningCartridges/<volser>**

Returns the same attributes as *GET /v1/cleaningCartridges* for a single cartridge.

Example 6-10 GET /v1/cleaningCartridges/<volser> output

```

[
  {
    "volser": "CLN041L1",
    "state": "importing",
    "accessible": "normal",
    "cleansRemaining": 50,
    "location": "slot_F1C3R5T1",
    "mediaType": "LTO",
    "mostRecentUsage": "2023-06-13T13:27:03-0700",
    "internalAddress": "FF030D"
  }
]

```


Data Cartridges

The following commands are for retrieving data about data cartridges:

► **GET /v1/dataCartridges**

Returns a list of data cartridges physically present in the library, as well as their physical location and other details.

Example 6-11 GET /v1/dataCartridges output

```
[
  {
    "volser": "0059RHL3",
    "state": "normal",
    "accessible": "normal",
    "location": "drive_F1C1R4",
    "mediaType": "LTO",
    "encrypted": "no",
    "mostRecentVerification": null,
    "mostRecentUsage": "2023-08-18T15:13:43-0700",
    "logicalLibrary": "LL1",
    "elementAddress": 257,
    "internalAddress": "010101"
  },
  {
    "volser": "0470RHL3",
    "state": "normal",
    "accessible": "normal",
    "location": "slot_F1C1R13T1",
    "mediaType": "LTO",
    "encrypted": "no",
    "mostRecentVerification": null,
    "mostRecentUsage": "2023-08-18T15:01:41-0700",
    "logicalLibrary": "LL1",
    "elementAddress": 1044,
    "internalAddress": "010414"
  },
  {
    "volser": "0617RHL3",
    "state": "normal",
    "accessible": "normal",
    "location": "slot_F1C3R33T1",
    "mediaType": "LTO",
    "encrypted": "no",
    "mostRecentVerification": null,
    "mostRecentUsage": "2023-08-18T13:22:14-0700",
    "logicalLibrary": "LL1",
    "elementAddress": 1034,
    "internalAddress": "01040A"
  },
  ...
]
```

Table 6-7 GET /v1/dataCartridges details

Attributes	Details
volser (string)	Volume serial number or bar code for the cartridge.
state (string)	Health status of the cartridge. Potential values are <i>normal</i> , <i>importing</i> , <i>exportQueued</i> , <i>uncertainBarcode</i> , <i>failedVerification</i> , <i>assignmentRequired</i> , <i>verifying</i> , and <i>unknown</i> .
accessible (string)	Indicates whether the accessor can reach the cartridge.
location (string)	Current location of the cartridge.
mediaType (string)	Indicates whether this is an LTO or 3592 cartridge.
encrypted (string)	Indicates whether the cartridge is encrypted (<i>yes</i>) or not (<i>no</i>). If the cartridge has not been mounted, the value is <i>null</i> .
mostRecentVerification (string)	Timestamp of the last time media verification was performed on this cartridge, or <i>null</i> if it has never been verified.
mostRecentUsage (string)	The last time this cartridge was mounted to a drive or <i>null</i> if it has never been mounted.
logicalLibrary (string)	The name of the logical library this cartridge is assigned to, or <i>null</i> if it is not assigned.
elementAddress (number)	The SCSI element address of this cartridge reported to the host. If it is not yet reported to the host, this value is <i>null</i> .
internalAddress (string)	A unique identifier for the cartridge, used to differentiate cartridges in the event of duplicate volsers.

► **GET /v1/dataCartridges/<volser>**

Returns the same attributes as *GET /v1/dataCartridges* for a specific cartridge.

Diagnostic Cartridges

The following commands are for retrieving information about diagnostic cartridges:

► GET /v1/diagnosticCartridges

Returns a list of diagnostic cartridges and their details from the library.

Example 6-12 GET /v1/diagnosticCartridges output

```
[
  {
    "volser": "DG I01L5",
    "state": "normal",
    "accessible": "normal",
    "location": "slot_F1C1R4T1",
    "mediaType": "LTO",
    "mostRecentUsage": "2023-05-24T16:24:09-0700",
    "internalAddress": "FF0403"
  },
  {
    "volser": "DG L08L9",
    "state": "normal",
    "accessible": "normal",
    "location": "slot_F1C1R37T1",
    "mediaType": "LTO",
    "mostRecentUsage": "2023-08-08T09:35:27-0700",
    "internalAddress": "FF0400"
  }
]
```

Table 6-8 GET /v1/diagnosticCartridges details

Attributes	Details
volser (string)	Volume serial number or bar code for the cartridge.
state (string)	Health status of the cartridge. Potential values are <i>normal</i> , <i>importing</i> , <i>exportQueued</i> , and <i>unknown</i> .
accessible (string)	Indicates whether the accessor can reach the cartridge.
location (string)	Current location of the cartridge.
mediaType (string)	Indicates whether this is an LTO or 3592 cartridge.
mostRecentUsage (string)	The last time this cartridge was mounted to a drive or <i>null</i> if it has never been mounted.
internalAddress (string)	A unique identifier for the cartridge, used to differentiate cartridges in the event of duplicate volsers.

► GET /v1/diagnosticCartridges<volser>

Returns the same attributes as *GET /v1/diagnosticCartridges* for a specific cartridge.

Drives

The following commands are about using tape drives:

► GET /v1/drives

Returns a list of tape drives and their details from the library.

Example 6-13 GET /v1/drives output

```
[
  {
    "location": "drive_F1C1R4",
    "sn": "63001A0703",
    "mediaType": "LTO",
    "state": "online",
    "operation": "ready",
    "accessible": "normal",
    "mtm": "3588-F9S",
    "barcode": "11S02XW367Y910WT002238",
    "interface": "fibreChannel",
    "interfaceMode": "single-mode",
    "logicalLibrary": "LL1",
    "use": "controlPath",
    "firmware": "N550",
    "encryption": "disabled",
    "wwnn": "50050763001a0703",
    "elementAddress": 257,
    "beacon": "disabled",
    "volser": "0059RHL3",
    "lastCleaned": null
  },
  {
    "location": "drive_F1C1R5",
    "sn": "63001A0704",
    "mediaType": "LTO",
    "state": "online",
    "operation": "empty",
    "accessible": "normal",
    "mtm": "3588-F8C",
    "barcode": "11S02XV407Y810WT020550",
    "interface": "fibreChannel",
    "interfaceMode": "multi-mode",
    "logicalLibrary": "LL1",
    "use": "controlPath",
    "firmware": "MA70",
    "encryption": "disabled",
    "wwnn": "50050763001a0704",
    "elementAddress": 258,
    "beacon": "disabled",
    "volser": null,
    "lastCleaned": null
  },
  ...
]
```

Table 6-9 GET /v1/drives details

Attributes	Details
location (string)	The location of the drive in the library, used for identification.
sn (string)	The drive serial number that is used for identification.
mediaType (string)	The media type supported by the drive, either <i>LTO</i> or <i>3592</i> .
state (string)	The health status of the tape drive. Potential values include <i>online</i> , <i>cleaning</i> , <i>updating</i> , <i>reset required</i> , <i>unreachable</i> , <i>initializing</i> , <i>restarting</i> , <i>inServiceMode</i> , and <i>unknown</i> .
operation (string)	The current operation that the tape drive is performing. Potential values include <i>empty</i> , <i>loading</i> , <i>unloading</i> , <i>ready</i> , <i>unloaded</i> , and <i>null</i> .
accessible (string)	Indicates whether the drive is reachable by the accessor.
mtm (string)	The machine type and model of the drive.
barcode (string)	The 11S bar code found on the drive.
interface (string)	The type of ports this drive includes. This is either <i>fibreChannel</i> or <i>SAS</i> .
interfaceMode (string)	The mode of the ports on the drive. The modes are <i>multi-mode</i> for FC-attached drives and <i>null</i> for SAS drives.
logicalLibrary (string)	The name of the logical library this drive is assigned to or <i>null</i> if it is not assigned.
use (string)	The assigned use of the drive in the logical library. This is either <i>access</i> , <i>controlPath</i> , or <i>verification</i> .
firmware (string)	The current firmware level of the tape drive.
encryption (string)	The state of encryption on the drive. Either <i>enabled</i> or <i>disabled</i> .
wwnn (string)	The worldwide node name of the drive.
elementAddress (number)	The SCSI element address of the drive reported to the host.
beacon (string)	The state of the beacon LED on the drive. This is either <i>enabled</i> or <i>disabled</i> .
volser (string)	The VOLSER of the cartridge that is mounted in the drive or <i>null</i> if no cartridge is present.
lastCleaned (string)	The timestamp of the last time the drive was cleaned.

- ▶ **GET /v1/drives/<location>**
Returns the same attributes as *GET /v1/drives* for a specific drive only.
- ▶ **GET /v1/drives<serial number>**
Returns the same attributes as *GET /v1/drives* for a specific drive only.
- ▶ **PUT /v1/drives/<location> {"beacon": <"enable" | "disable">}**
Sets the LED beacon on the back of the drive. Either on (*enable*) or off (*disable*).
- ▶ **PUT /v1/drives/<location> {"use": <"access" | "controlPath" | "verification">}**
Changes the usage of a drive in a logical library. It can be set as a normal data access drive, control path drive for communicating with the host, or a media verification drive.
- ▶ **POST /v1/drives/<location>/clean**
Moves a cleaning cartridge to the drive in the specified location and begins cleaning it.
- ▶ **POST /v1/drives/<location>/reset {"mode": <"normal" | "hard" >}**
Restart a drive. Normal mode sends a command to the drive to perform a soft reset. Hard mode physically power cycles the drive. A hard reset should be performed only when it is known that there is no media loaded in the drive. Use caution when running this command.

Ethernet Ports

The following commands are about using Ethernet ports:

- ▶ **GET /v1/ethernetPorts**
Returns a list of Ethernet ports and their settings from the library.

Example 6-14 GET /v1/ethernetPorts output

```
[
  {
    "location": "ethernetPort_F1Pa",
    "macAddress": "40:F2:E9:52:D1:C1",
    "ipv4Address": "9.11.44.28",
    "ipv4Subnet": "255.255.252.0",
    "ipv4Gateway": "9.11.44.1",
    "ipv4Assignment": "static",
    "ipv4Primary": null,
    "ipv4Secondary": null,
    "ipv6Address": "disabled",
    "ipv6PrefixLength": "disabled",
    "ipv6Gateway": "disabled",
    "ipv6Primary": "disabled",
    "ipv6Secondary": "disabled",
    "ipv6Link": "disabled",
    "ipv6DHCP": "disabled",
    "ipv6StatelessConfig": "disabled",
    "ipv6Static": "disabled"
  },
  {
    "location": "ethernetPort_F1Pb",
    "macAddress": "40:F2:E9:52:D1:C2",
    "ipv4Address": "disabled",
    "ipv4Subnet": "disabled",
    "ipv4Gateway": "disabled",
```

```

    "ipv4Assignment": "disabled",
    "ipv4Primary": "disabled",
    "ipv4Secondary": "disabled",
    "ipv6Address": "disabled",
    "ipv6PrefixLength": "disabled",
    "ipv6Gateway": "disabled",
    "ipv6Primary": "disabled",
    "ipv6Secondary": "disabled",
    "ipv6Link": "disabled",
    "ipv6DHCP": "disabled",
    "ipv6StatelessConfig": "disabled",
    "ipv6Static": "disabled"
  },
  {
    "location": "ethernetPort_F1Pimc",
    "macAddress": "40:F2:E9:52:D1:BF",
    "ipv4Address": "192.168.2.2",
    "ipv4Subnet": null,
    "ipv4Gateway": null,
    "ipv4Assignment": "static",
    "ipv4Primary": null,
    "ipv4Secondary": null,
    "ipv6Address": "FDC0:FFEE:BAFE::251:1",
    "ipv6PrefixLength": "0",
    "ipv6Gateway": null,
    "ipv6Primary": null,
    "ipv6Secondary": null,
    "ipv6Link": "FE80::42F2:E9FF:FE52:D1BF",
    "ipv6DHCP": "disabled",
    "ipv6StatelessConfig": "disabled",
    "ipv6Static": "enabled"
  },
  {
    "location": "ethernetPort_F1Ptssc",
    "macAddress": "40:F2:E9:52:D1:C0",
    "ipv4Address": "disabled",
    "ipv4Subnet": "disabled",
    "ipv4Gateway": "disabled",
    "ipv4Assignment": "disabled",
    "ipv4Primary": "disabled",
    "ipv4Secondary": "disabled",
    "ipv6Address": "disabled",
    "ipv6PrefixLength": "disabled",
    "ipv6Gateway": "disabled",
    "ipv6Primary": "disabled",
    "ipv6Secondary": "disabled",
    "ipv6Link": "disabled",
    "ipv6DHCP": "disabled",
    "ipv6StatelessConfig": "disabled",
    "ipv6Static": "disabled"
  },
  {
    "location": "ethernetPort_F1Pservice",
    "macAddress": "40:F2:E9:52:D1:BF",
    "ipv4Address": "192.168.1.2",

```

```

    "ipv4Subnet": null,
    "ipv4Gateway": null,
    "ipv4Assignment": "static",
    "ipv4Primary": null,
    "ipv4Secondary": null,
    "ipv6Address": "FDC0:FFEE:BABE::250:1",
    "ipv6PrefixLength": "0",
    "ipv6Gateway": null,
    "ipv6Primary": null,
    "ipv6Secondary": null,
    "ipv6Link": "FE80::42F2:E9FF:FE52:D1BF",
    "ipv6DHCP": "disabled",
    "ipv6StatelessConfig": "disabled",
    "ipv6Static": "enabled"
  }
]

```

Table 6-10 GET /v1/ethernetPorts details

Attributes	Details
location (string)	Location of the Ethernet port that is used for identification.
macAddress (string)	The MAC address of the port.
ipv4Address (string)	The IPv4 address of the port.
ipv4Subnet (string)	The IPv4 subnet mask of the port.
ipv4Gateway (string)	The IPv4 gateway address of the port.
ipv4Assignment (string)	Indicates a <i>static</i> or <i>dynamic</i> IP address assignment.
ipv4Primary (string)	The primary IPv4 DNS server for the port.
ipv4Secondary (string)	The secondary IPv4 DNS server for the port.
ipv6Address (string)	The IPv6 address of the port.
ipv6PrefixLength (string)	The IPv6 prefix length of the port.
ipv6Gateway (string)	The IPv6 gateway address of the port.
ipv6Primary (string)	The primary IPv6 DNS server for the port.
ipv6Secondary (string)	The secondary IPv6 DNS server for the port.
ipv6Link (string)	The IPv6 link setting for the port.
ipv6DHCP (string)	The IPv6 DHCP setting for the port.
ipv6StatelessConfig (string)	The IPv6 stateless config setting for the port.
ipv6Static (string)	The IPv6 static IP setting for the port.

► **GET /v1/ethernetPorts/<location>**

Returns the same attributes as *GET /v1/ethernetPorts* for a single port.

- ▶ **PUT /v1/ethernetPorts/<location> {"ipv4Address": <IPv4 address>, "ipv4Subnet": <IPv4 subnet mask>, "ipv4Gateway": <IPv4 gateway address>, "ipv4Assignment": <"static"|"dynamic">, "ipv4Primary": <IPv4 address>, "ipv4Secondary": <IPv4 address>}**

Change network settings on a given Ethernet port. Warning: this may result in loss of network access if using with REST over Ethernet on the port that is currently connected.

Events:

- ▶ **GET /v1/events**

Returns a list of error, warning, and informational events and details about them from the library.

Example 6-15 GET /v1/events output

```
[
  {
    "ID": 4885,
    "severity": "information",
    "time": "2023-08-21T11:44:19-0700",
    "location": null,
    "user": "SCSIControlPath_F1C2R2",
    "description": "An administrator password was changed",
    "state": null,
    "errorCode": "0810"
  },
  {
    "ID": 4884,
    "severity": "information",
    "time": "2023-08-21T11:43:49-0700",
    "location": null,
    "user": "SCSIControlPath_F1C2R2",
    "description": "An administrator password was changed",
    "state": null,
    "errorCode": "0810"
  },
  {
    "ID": 4883,
    "severity": "information",
    "time": "2023-08-21T11:43:32-0700",
    "location": null,
    "user": "SCSIControlPath_F1C2R2",
    "description": "An administrator password was changed",
    "state": null,
    "errorCode": "0810"
  },
  {
    "ID": 4882,
    "severity": "information",
    "time": "2023-08-21T11:38:13-0700",
    "location": null,
    "user": "admin",
    "description": "A user logged in to the GUI from 9.61.250.239",
    "state": null,
    "errorCode": "0800"
  },
]
```

```

{
  "ID": 4881,
  "severity": "information",
  "time": "2023-08-21T11:34:57-0700",
  "location": null,
  "user": "admin",
  "description": "A user logged out of the GUI from 9.11.45.65",
  "state": null,
  "errorCode": "0812"
},
{
  "ID": 4880,
  "severity": "information",
  "time": "2023-08-21T11:34:19-0700",
  "location": null,
  "user": "admin",
  "description": "A user logged in to the GUI from 9.11.45.65",
  "state": null,
  "errorCode": "0800"
},
...
]

```

Table 6-11 GET /v1/events details

Attribute	Details
ID (number)	The unique ID of the event on the library.
severity (string)	Severity of the event. Values include <i>error</i> , <i>warning</i> , <i>inactive error</i> , <i>inactive warning</i> , and <i>informational</i> .
time (string)	The time when the event occurred.
location (string)	The hardware resource where the event occurred. This can be <i>null</i> if no hardware was involved.
user (string)	The ID of the user that issued the action that triggered the event. This can also be <i>System</i> if there was no user involved.
description (string)	A description of what specifically went wrong.
state (string)	The current state of the event regarding its call home status. This can be <i>null</i> if no call home has occurred yet (or will not occur).
errorCode (string)	A four-digit hex value used to uniquely identify error types in the library.

► **GET /v1/events/<ID>**

Returns the same attributes as *GET /v1/events* for a specific event.

► **GET /v1/events?after=<timestamp> and GET /v1/events?before=<timestamp>**

Returns the same attributes as *GET /v1/events* for any events that have occurred before or after the given time stamp.

► **POST /v1/events/<ID>/active**

Set an inactive event to active for reporting. This will respond with an HTTP error code of 201 if it was successfully received by the library.

► **POST /v1/events/<ID>/inactive**

Set an active event to inactive after the problem has been resolved. This will respond with an HTTP error code of 201 if it was successfully received by the library.

FC Ports

The following commands describe some commands for FC ports:

► **GET /v1/fcPorts**

Returns a list of fibre channel ports and their details from the library.

Example 6-16 GET /v1/fcPorts output

```
[
  {
    "location": "fcPort_F1C1R4P0",
    "driveLocation": "drive_F1C1R4",
    "driveSn": "63001A0703",
    "portNumber": 0,
    "state": "noLightDetected",
    "wwpn": "50050763005a0703",
    "speedSetting": "auto",
    "speedActual": "1Gbps",
    "topologySetting": "auto-L",
    "topologyActual": "unknown",
    "loopID": 4
  },
  {
    "location": "fcPort_F1C1R4P1",
    "driveLocation": "drive_F1C1R4",
    "driveSn": "63001A0703",
    "portNumber": 1,
    "state": "noLightDetected",
    "wwpn": "50050763009a0703",
    "speedSetting": "auto",
    "speedActual": "unknown",
    "topologySetting": "auto-L",
    "topologyActual": "unknown",
    "loopID": 68
  },
  {
    "location": "fcPort_F1C1R5P0",
    "driveLocation": "drive_F1C1R5",
    "driveSn": "63001A0704",
    "portNumber": 0,
    "state": "noLightDetected",
    "wwpn": "50050763005a0704",
    "speedSetting": "auto",
    "speedActual": "1Gbps",
    "topologySetting": "auto-L",
    "topologyActual": "unknown",
    "loopID": 5
  }
]
```

```

    },
    {
      "location": "fcPort_F1C1R5P1",
      "driveLocation": "drive_F1C1R5",
      "driveSn": "63001A0704",
      "portNumber": 1,
      "state": "communicationEstablished",
      "wwpn": "50050763009a0704",
      "speedSetting": "auto",
      "speedActual": "4Gbps",
      "topologySetting": "auto-L",
      "topologyActual": "L-Port",
      "loopID": 69
    },
    {
      "location": "fcPort_F1C3R3P0",
      "driveLocation": "drive_F1C3R3",
      "driveSn": "63001A070C",
      "portNumber": 0,
      "state": "noLightDetected",
      "wwpn": "50050763005a070c",
      "speedSetting": "auto",
      "speedActual": "1Gbps",
      "topologySetting": "auto-L",
      "topologyActual": "unknown",
      "loopID": 11
    },
    {
      "location": "fcPort_F1C3R3P1",
      "driveLocation": "drive_F1C3R3",
      "driveSn": "63001A070C",
      "portNumber": 1,
      "state": "noLightDetected",
      "wwpn": "50050763009a070c",
      "speedSetting": "auto",
      "speedActual": "unknown",
      "topologySetting": "auto-L",
      "topologyActual": "unknown",
      "loopID": 75
    }
  ]

```

Table 6-12 GET /v1/fcPorts details

Attribute	Details
location (string)	The location of the port.
driveLocation (string)	The location of the drive that the port is attached to.
driveSn (string)	The serial number of the drive that the port is attached to.
portNumber (number)	The port number on the drive, either 0 or 1.

Attribute	Details
state (string)	The state of the port. Describes whether it is connected or not.
wwpn (string)	The worldwide port name of the port.
speedSetting (string)	The setting for the port's data transfer speed.
speedActual (string)	The data transfer speed that the port is currently operating at.
topologySetting (string)	The current setting for the topology or protocol of the port.
topologyActual (string)	The topology or protocol that the port is currently using.
loopID (string)	A unique value that identifies the position of the drive in the library.

► **GET /v1/fcPorts/<location>**

Returns the same attributes as *GET /v1/fcPorts* for a single port.

Frames

The following commands describe retrieving information about frames:

► **GET /v1/frames**

Returns a list of frames installed in a library along with their details. A Diamondback library can have only one frame, so the list always contains one entry.

Example 6-17 GET /v1/frames output

```
[
  {
    "location": "frame_F1",
    "state": "normal",
    "type": "base",
    "mtm": "4884-L9A",
    "sn": "13B0017",
    "mediaType": "LTO",
    "frontDoor": "closed",
    "frontDoorLastChanged": "2023-08-16T16:28:12-0700",
    "rearDoor": "open",
    "rearDoorLastChanged": "2023-08-16T10:01:33-0700",
    "sideDoor": null,
    "sideDoorLastChanged": null,
    "slots": 1584,
    "cartridges": 68,
    "drives": 5,
    "ioStations": 2
  }
]
```

► **GET /v1/frames/<location>**

Returns the same attributes as *GET /v1/frames* for a single frame. Because a Diamondback library is a single frame, this will always be the same data returned by *GET /v1/frames*.

GUI Settings

Use the **GET /v1/guiSettings** command to view limited settings from the Diamondback web GUI.

Example 6-18 GET /v1/guiSettings output

```
[
  {
    "autoLogout": 30,
    "autoIMCLogin": null
  }
]
```

Table 6-13 GET /v1/guiSettings details

Attribute	Details
autoLogout (number)	The number of minutes until a user is logged out of a web session.
autoIMCLogin (string)	Not applicable to Diamondback libraries. This value will always be <i>null</i> .

I/O Stations

Use the following commands to retrieve information about I/O stations:

► **GET /v1/ioStations**

Returns details about the I/O station installed in the library.

Example 6-19 GET /v1/ioStations output

```
[
  {
    "location": "ioStation_C3",
    "state": "normal",
    "door": null,
    "magazine": {
      "mediaType": "LT0",
      "ioSlots": 10,
      "contentsVolser": [
        null,
        null,
        null,
        null,
        null,
        null,
        null,
        null,
        null,
        null
      ]
    }
  }
]
```

```

    "contentsInternalAddress": [
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null
    ]
  }
}

```

Table 6-14 GET /v1/ioStations details

Attributes	Details
location (string)	The location of the I/O station in the library.
state (string)	The status of the I/O station. The only possible values for Diamondback libraries are <i>normal</i> and <i>unknown</i> . If the I/O magazine is removed, it will not show up in the request results.
door (string)	Not applicable to Diamondback libraries as there is no I/O station door. This value will always be <i>null</i> .
magazine (object)	A JSON object describing the current state of the I/O magazine.

The magazine attribute contains a nested structure with the following attributes:

Table 6-15 GET /v1/ioStations magazine attributes details

Attributes	Details
mediaType (string)	The media type supported by the magazine, either <i>LTO</i> or <i>3592</i> .
ioSlots (number)	The number of slots in the magazine.
contentsVolser (array of strings)	A list of the VOLSERs of cartridges currently in the magazine.
contentsInternalAddress (array of strings)	A list of cartridge internal addresses currently in the magazine.

► **GET /v1/ioStations/<location>**

Returns the same attributes as *GET /v1/frames* for a single frame. Because a Diamondback library is a single frame, this will always be the same data returned by *GET /v1/frames*.

Figure 6-3

Library

Issue the following commands to use the tape library:

► **GET /v1/library**

Returns general information about the library configuration.

Example 6-20 GET /v1/library output

```
[
  {
    "name": "Ravioli",
    "status": "driveDegraded",
    "cartridgeAccess": "normal",
    "totalCapacity": 1584,
    "licensedCapacity": 800,
    "totalCartridges": 68,
    "assignedCartridges": 56,
    "firmware": "2.9.0.3-11K",
    "sn": "13B0017",
    "time": "2023-08-21T12:07:19-0700",
    "timezone": "America/Phoenix",
    "ntpMode": "disabled",
    "primaryNtpAddress": null,
    "secondaryNtpAddress": null,
    "location": null,
    "address": null,
    "city": null,
    "state": null,
    "country": null,
    "contact": null,
    "telephone": null,
    "secondaryTelephone": null,
    "secureCommunications": "enabled",
    "autoEjectCleaningCartridges": "enabled",
    "elasticCapacity": "maxCapacity",
    "activeAccessors": "accessorAOnly",
    "vioStatus": "enabled",
    "nmaDetection": "enabled",
    "capacityUtilThresh": 98.0,
    "dualAccessorUtilThresh": 98.0
  }
]
```

Table 6-16 GET /v1/library details

Attribute	Details
name (string)	The name of the library.
status (string)	The health status of the library. The default good status of the library is <i>online</i> .
cartridgeAccess (string)	Reports whether the accessor can access cartridges or not.

Attribute	Details
totalCapacity (number)	Total physical capacity for tape cartridges in the library.
licensedCapacity (number)	Total capacity for tape cartridges licensed for the library.
totalCartridges (number)	Total cartridges currently present in the library.
assignedCartridges (number)	Total cartridges currently assigned to logical libraries.
firmware (string)	The firmware level of the library.
sn (string)	The library's serial number.
time (string)	The current time according to the library.
timezone (string)	The current timezone of the library.
ntpMode (string)	Whether or not NTP is enabled on the library. Values are either <i>enabled</i> or <i>disabled</i> .
primaryNtpAddress (string)	The IP address of the primary NTP server, or <i>null</i> if no NTP servers are configured.
secondaryNtpAddress (string)	The IP address of the secondary NTP server, or <i>null</i> if it is not configured.
location (string)	The user-entered physical location of the library, or <i>null</i> if it is not configured.
address (string)	The user-entered physical address of the library, or <i>null</i> if it is not configured.
city (string)	The user-entered city in which the library is located, or <i>null</i> if it is not configured.
state (string)	The user-entered state in which the library is located, or <i>null</i> if it is not configured.
country (string)	The user-entered country in which the library is located, or <i>null</i> if it is not configured.
contact (string)	The name of the primary contact for this library, used by IBM support. This value is <i>null</i> if not configured.
telephone (string)	The primary contact telephone number for this library, used by IBM support. This value is <i>null</i> if not configured.
secondaryTelephone (string)	The secondary telephone number for the primary contact for this library, used by IBM support. This value is <i>null</i> if not configured.

Attribute	Details
secureCommunications (string)	Indicates whether the library requires an HTTPS connection for the web GUI and REST over Ethernet. This value cannot be changed on Diamondback libraries and is always set to <i>enabled</i> .
autoEjectCleaningCartridges (string)	The current setting for automatically ejecting cleaning cartridges to the I/O station when they have no cleans remaining. The values are either <i>enabled</i> or <i>disabled</i> .
elasticCapacity (string)	Not applicable to Diamondback libraries. This value is always <i>null</i> .
activeAccessors (string)	Not applicable to Diamondback libraries. This value is always <i>null</i> .
vioStatus (string)	The current setting for Virtual I/O. This value is always <i>enabled</i> on Diamondback libraries.
nmaDetection (string)	The current setting for NMA detection. This value is always <i>enabled</i> on Diamondback libraries.
capacityUtilThresh (number)	A percentage indicating how full the library must be before triggering warning events.
dualAccessorUtilThresh (number)	Not applicable to Diamondback libraries as there will only ever be one robotic accessor.

► **GET /v1/library/saveConfig**

This endpoint does not have any attributes. It simply returns an HTTP header followed by a Diamondback configuration file that can be used to restore library configuration if needed.

► **PUT /v1/library/name {"name": <library name>}**

Change the library name to the string specified in the JSON data. A successful update will report a return code 201.

► **PUT /v1/library/time {"time": <time>}**

Change the system time of the library to the time specified in the JSON data. Time string format is "YYYY-MM-DDThh:mm:ss±hh:mm". A successful update reports a return code 201.

► **PUT /v1/library/timezone {"timezone": <timezone>}**

Set the library time zone. Time zone formats are specified by the Internet Assigned Numbers Authority (IANA). For a list of time zones, see the [IANA Time Zone Database](#). A successful update reports a return code 201.

► **POST /v1/library/reset**

Performs a soft library restart. This should not impact host operations. However, the command might not return a status to the host because the library with which it is communicating with has been reset. A successful update reports a return code 201.

► **POST /v1/library/restoreConfig**

Performs a configuration restore on the library from a given configuration backup file. A successful update reports a return code 201.

► **POST /v1/library/prepareForShip**

Moves all cartridges to the proper location in order to prepare a library to be moved to a different location. A successful update reports a return code 201.

Logical libraries

Use the following commands to list information about logical libraries:

► **GET /v1/logicalLibraries**

Returns a list of logical libraries and their details from the library.

Example 6-21 GET /v1/logicalLibraries output

```
[
  {
    "name": "LL1",
    "mediaType": "LTO",
    "drives": 4,
    "virtualSlots": 1584,
    "virtualIOSlots": 255,
    "cartridges": 56,
    "encryptionMethod": "none"
  }
]
```

Table 6-17 Get /v1/logicalLibraries details

Attributes	Details
name (string)	The name of the logical library.
mediaType (string)	The media type of the logical library (<i>LTO</i> or <i>3592</i>).
drives (number)	The number of drives in the logical library.
virtualSlots (number)	The number of virtual storage slots reported to the host.
virtualIOSlots (number)	The number of virtual I/O slots reported to the host.
cartridges (number)	The number of cartridges assigned to the logical library.
encryptionMethod (string)	The encryption method that is currently enabled for the logical library. Possible values include <i>none</i> , <i>systemManaged</i> , <i>applicationManaged</i> , and <i>libraryManagedBarcode</i> .

► **GET /v1/logicalLibraries/<name>**

Returns the same attributes as *GET /v1/logicalLibraries* for a single logical library.

► **GET /v1/logicalLibraries/<name>/volserRanges**

Returns a list of VOLSER ranges assigned to each logical library.

Example 6-22 GET /v1/logicalLibraries/<name>/volserRanges

```
[
  {
```

```

    "start": "000000",
    "end": "ZZZZZ",
    "cartridges": 65
  }
]

```

Example 6-23 GET /v1/logicalLibraries/<name>/volserRanges

Attributes	Details
start (string)	The first VOLSER in the range.
end (string)	The last VOLSER in the range.
cartridges (number)	The number of cartridges assigned to the logical library in this VOLSER range.

Logs

Use the following commands to view and collect logs:

► GET /v1/logs

Returns a list of log files that can be exported from the library.

Example 6-24 GET /v1/logs output

```

[
  {
    "filename": "",
    "state": "InProgress",
    "startTime": "",
    "lastUpdateTime": "",
    "percentComplete": 50
  },
  {
    "filename": "Diamondback_SS_081E_B0017_20230821112129.zip",
    "state": "completed",
    "startTime": "2023-08-21T11:21:29-0700",
    "lastUpdateTime": "2023-08-21T11:27:09-0700",
    "percentComplete": 100
  },
  {
    "filename": "Diamondback_SS_9180_B0017_20230818160633.zip",
    "state": "completed",
    "startTime": "2023-08-18T16:06:33-0700",
    "lastUpdateTime": "2023-08-18T16:08:35-0700",
    "percentComplete": 100
  },
  {
    "filename": "Diamondback_LOG_B0017_20230818150229.zip",
    "state": "completed",
    "startTime": "2023-08-18T15:02:29-0700",
    "lastUpdateTime": "2023-08-18T15:04:08-0700",
    "percentComplete": 100
  },
  {
    "filename": "Diamondback_LOG_B0017_20230818132317.zip",

```

```

    "state": "completed",
    "startTime": "2023-08-18T13:23:17-0700",
    "lastUpdateTime": "2023-08-18T13:25:50-0700",
    "percentComplete": 100
  },
  {
    "filename": "Diamondback_LOG_B0017_20230818123845.zip",
    "state": "completed",
    "startTime": "2023-08-18T12:38:45-0700",
    "lastUpdateTime": "2023-08-18T12:41:21-0700",
    "percentComplete": 100
  },
  {
    "filename": "Diamondback_LOG_B0017_20230816154300.zip",
    "state": "completed",
    "startTime": "2023-08-16T15:43:00-0700",
    "lastUpdateTime": "2023-08-16T15:45:37-0700",
    "percentComplete": 100
  }
]

```

Table 6-18 GET /v1/logs details

Attributes	Details
filename (string)	The unique filename of the log file.
state (string)	The current state of the log collection associated with the file. Possible values are <i>inProgress</i> , <i>completed</i> , and <i>failed</i> .
startTime (string)	The time when the log collection associated with the file started.
lastUpdateTime (string)	The time when the log file was last updated.
percentComplete (number)	A percentage indicating how far along the log collection process is for the file.

► **GET /v1/logs/<filename>**

Returns the same attributes as *GET /v1/logs* for a single log file.

► **GET /v1/logs/<filename>/export**

This endpoint is used for exporting a specific log file. This request returns an HTTP header followed by the requested file.

► **POST /v1/logs {"location": "library"}**

Begins a library log collection process. When this is complete, the logs can be retrieved by using the *GET /v1/logs/<filename>/export* command.

Returns an HTTP header indicating success, followed by a JSON structure with the name of the log file that is collected. This can be queried with *GET /v1/logs/<filename>* to monitor the process.

Example 6-25 POST /v1/logs{"location": "library"}

```
{
```

```
    "filename": "Diamondback_LOG_B0017_20230821121302.zip"
  }
}
```

► **POST /v1/logs {"location": "drive_F<frame>C<column>R<row>"}**

Begins a drive log collection process. When this is complete, the logs can be retrieved by using the `GET /v1/logs/<filename>/export` command.

Example 6-26 POST /v1/logs {"location": "drive_F<frame>C<column>R<row>"} output

```
{
  "filename": "Diamondback_DriveLogs_F1C3R3_20230821121421.zip"
}
```

Node cards

Use the following commands to list information and reset node cards:

► **GET /v1/nodeCards**

Returns a list of node cards that are installed in the library and their details.

Example 6-27 GET /v1/nodeCards output

```
[
  {
    "ID": 65,
    "type": "LCC",
    "location": "frame_F1",
    "state": "online",
    "partNum": "03FR850",
    "sn": "Y010MY2AK20D",
    "barcode": "11S03FR850Y010MY2AK20D",
    "ec": 0,
    "firmware": "2.9.0.3-11K",
    "cfBarcode": "SPG2021102100009",
    "cfPartNum": "SG9CF16GHYDCIBM",
    "cfVendor": "SMART",
    "reportingLCC": "yes",
    "primaryLCC": "yes",
    "lastRestart": "2023-08-21T11:23:16-0700"
  },
  {
    "ID": 129,
    "type": "RCC",
    "location": "accessor_Aa",
    "state": "online",
    "partNum": "03FR700",
    "sn": "YN10C0000054",
    "barcode": "11S03FR700YN10C0000054",
    "ec": 0,
    "firmware": "2100",
    "cfBarcode": null,
    "cfPartNum": null,
    "cfVendor": null,
    "reportingLCC": null,
    "primaryLCC": null,
    "lastRestart": null
  }
]
```

```
}
]
```

Table 6-19 GET /v1/nodeCards details

Attributes	Details
ID (number)	The unique ID of the node card in the library.
type (string)	The type of the node card. Possible values are <i>LCC</i> and <i>RCC</i> .
location (string)	The location of the node card in the library.
state (string)	The current state of the node card.
partNum (string)	The part number of the node card.
sn (string)	The serial number of the node card.
barcode (string)	The bar code found on the node card.
ec (number)	The EC number of the node card.
firmware (string)	The current firmware level of the node card.
cfBarcode (string)	The bar code on the LCC's CompactFlash card. If this is not an LCC, the value is <i>null</i> .
cfPartNum (string)	The part number of the LCC's CompactFlash card. If this is not an LCC, the value is <i>null</i> .
cfVendor (string)	The vendor of the LCC's CompactFlash card. If this is not an LCC, the value is <i>null</i> .
reportingLCC (string)	Whether this is the LCC reporting information to the requester. If this node card is an LCC, the value is <i>yes</i> because there is only one LCC in a Diamondback. If this node card is not an LCC, the value is <i>null</i> .
primaryLCC (string)	Whether this is the primary LCC. If this node card is an LCC, the value is <i>yes</i> . If this node card is not an LCC, the value is <i>null</i> .
lastRestart (string)	The last time this node card was reset.

► **GET /v1/nodeCards/<ID>**

Returns the same attributes as *GET /v1/nodeCards* for a single node card in the library. *POST /v1/nodeCards/<ID>/reset*

Resets the node card with the specified ID. A successful update reports a return code 201.

► **POST /v1/nodeCards/<ID>/reset**

Resets the node card with the given ID. A successful update reports a return code 201.

Notification

Use the following commands to view syslog servers:

► **GET /v1/notification/syslog/servers**

Returns a list of syslog servers configured for the library and their details.

Example 6-28 GET /v1/notification/syslog/servers output

```
[
  {
    "address": "1.1.1.1",
    "port": 514,
    "subscribed": [
      "warning"
    ]
  }
]
```

Table 6-20 GET /v1/notification/syslog/servers details

Attributes	Details
address (string)	The IP address of the syslog server.
port (number)	The port number used by the syslog server.
subscribed (array of strings)	An array of strings for types of events to report over syslog. Possible values include <i>error</i> , <i>warning</i> , and <i>information</i> .

► **GET /v1/notification/syslog/servers/<address>**

Returns the same attributes as *GET /v1/notifications/syslog/servers* for a single syslog server.

Power Supplies

Use the following commands to view power supply information:

► **GET /v1/powerSupplies**

Returns a list of power supplies installed in the library and their current states.

Example 6-29 GET /v1/powerSupplies

```
[
  {
    "location": "powerSupply_F1PSa",
    "state": "failed"
  },
  {
    "location": "powerSupply_F1PSb",
    "state": "online"
  }
]
```


Table 6-21 GET /v1/powerSupplies details

Attributes	Details
location (string)	The location of the power supply.
state (string)	The current state of the power supply. Possible values are <i>online</i> , <i>failed</i> , and <i>unknown</i> .

► **GET /v1/powerSupplies/<location>**

Returns the same attributes as *GET /v1/powerSupplies* for a single power supply.

Reports

Use the following commands to view activity reports:

► **GET /v1/reports/accessors**

Returns activity reports for the robotic accessor installed in the library.

Example 6-30 GET /v1/reports/accessors output

```
[
  {
    "location": "accessor_Aa",
    "time": "2023-08-21T12:05:00-0700",
    "duration": 3600,
    "pivots": 0,
    "barCodeScans": 316,
    "travelX": 3,
    "travelY": 12,
    "getsGripper1": 0,
    "putsGripper1": 0,
    "getsGripper2": 0,
    "putsGripper2": 0,
    "temperatureAverage": 22.4,
    "temperatureMin": 22.3,
    "temperatureMax": 22.7,
    "humidityAverage": 64.5,
    "humidityMin": 63.6,
    "humidityMax": 65.4
  },
  {
    "location": "accessor_Aa",
    "time": "2023-08-21T11:05:00-0700",
    "duration": 3600,
    "pivots": 0,
    "barCodeScans": 0,
    "travelX": 0,
    "travelY": 0,
    "getsGripper1": 0,
    "putsGripper1": 0,
    "getsGripper2": 0,
    "putsGripper2": 0,
    "temperatureAverage": 22.2,
    "temperatureMin": 22.2,
    "temperatureMax": 22.3,
    "humidityAverage": 65.8,
  }
]
```

```

    "humidityMin": 65.6,
    "humidityMax": 66.2
  },
  {
    "location": "accessor_Aa",
    "time": "2023-08-21T10:05:00-0700",
    "duration": 3600,
    "pivots": 0,
    "barCodeScans": 0,
    "travelX": 0,
    "travelY": 0,
    "getsGripper1": 0,
    "putsGripper1": 0,
    "getsGripper2": 0,
    "putsGripper2": 0,
    "temperatureAverage": 22.2,
    "temperatureMin": 22.2,
    "temperatureMax": 22.3,
    "humidityAverage": 66.2,
    "humidityMin": 66.2,
    "humidityMax": 66.3
  },
  ...
]

```

Table 6-22 GET /v1/reports/accessors details

Attributes	Details
location (string)	The location of the robotic accessor.
time (string)	The time when the data was recorded.
duration (number)	The number of seconds over which this entry was recorded.
pivots (number)	Not applicable to Diamondback libraries. This value will always be zero.
barcodeScans (number)	The number of bar code scans performed over the duration.
travelX (number)	How far the accessor has moved horizontally (in meters) over the duration.
travelY (number)	How far the accessor has moved vertically (in meters) over the duration.
getsGripper1 (number)	The number of cartridge retrievals the first gripper has performed over the duration.
putsGripper1 (number)	The number of cartridge placements the first gripper has performed over the duration.
getsGripper2 (number)	The number of cartridge retrievals the second gripper has performed over the duration.

Attributes	Details
putsGripper2 (number)	The number of cartridge placements the second gripper has performed over the duration.
temperatureAverage (number)	The average temperature the accessor recorded over the duration (in Celsius), accurate to one decimal place.
temperatureMin (number)	The minimum temperature the accessor has recorded over the duration (in Celsius), accurate to one decimal place.
temperatureMax (number)	The maximum temperature the accessor has recorded over the duration (in Celsius), accurate to one decimal place.
humidityAverage (number)	The average humidity percentage the accessor has recorded over the duration, accurate to one decimal place.
humidityMin (number)	The minimum humidity percentage the accessor has recorded over the duration, accurate to one decimal place.
humidityMax (number)	The maximum humidity percentage the accessor has recorded over the duration, accurate to one decimal place.

► **GET /v1/reports/drives**

Returns activity reports for all tape drives installed in the library.

Example 6-31 GET /v1/reports/drives output

```
[
  {
    "location": "drive_F1C1R4",
    "sn": "63001A0703",
    "time": "2023-08-21T12:05:00-0700",
    "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": 68.0,
    "humidityMin": 68.0,
    "humidityMax": 68.0
  },
]
```

```

{
  "location": "drive_F1C1R5",
  "sn": "63001A0704",
  "time": "2023-08-21T12:05:00-0700",
  "duration": 3600,
  "mounts": 0,
  "cleans": 0,
  "dataReadByHosts": 0,
  "dataWrittenByHosts": 0,
  "dataWrittenToCartridges": 0,
  "errorsCorrectedRead": 0,
  "errorsCorrectedWrite": 0,
  "errorsUncorrected": 0,
  "temperatureAverage": 22.5,
  "temperatureMin": 22.0,
  "temperatureMax": 23.0,
  "humidityAverage": 0.0,
  "humidityMin": 0.0,
  "humidityMax": 0.0
},
{
  "location": "drive_F1C3R3",
  "sn": "63001A070C",
  "time": "2023-08-21T12:05:00-0700",
  "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": 71.0,
  "humidityMin": 71.0,
  "humidityMax": 71.0
},
...
]

```

Table 6-23 GET /v1/reports/drives details

Attributes	Details
location (string)	The location of the drive.
sn (string)	The serial number of the drive.
time (string)	The time at which the data was recorded.
duration (number)	The number of seconds over which this entry was recorded.

Attributes	Details
mounts (number)	The number of cartridges that were mounted to this drive over the duration.
cleans (number)	The number of times this drive was cleaned over the duration.
dataReadByHosts (number)	The amount of data (in MB) that was read from cartridges by this drive over the duration.
dataWrittenByHosts (number)	The amount of data (in MB) that was written to cartridges from the host by this drive over the duration.
dataWrittenToCartridges (number)	The amount of data (in MB) that was written to cartridges by this drive over the duration. The average compression ratio can be found by dividing <i>dataWrittenByHosts</i> by <i>dataWrittenToCartridges</i> .
errorsCorrectedRead (number)	The number of corrected read errors on this drive over the duration.
errorsCorrectedWrite (number)	The number of corrected write errors on this drive over the duration.
errorsUncorrected (number)	The number of uncorrected errors on this drive over the duration.
temperatureAverage (number)	The average temperature the drive has recorded over the duration (in Celsius), accurate to one decimal place.
temperatureMin (number)	The minimum temperature the drive has recorded over the duration (in Celsius), accurate to one decimal place.
temperatureMax (number)	The maximum temperature the drive has recorded over the duration (in Celsius), accurate to one decimal place.
humidityAverage (number)	The average humidity percentage the drive has recorded over the duration, accurate to one decimal place.
humidityMin (number)	The minimum humidity percentage the drive has recorded over the duration, accurate to one decimal place.
humidityMax (number)	The maximum humidity percentage the drive has recorded over the duration, accurate to one decimal place.

► **GET /v1/reports/library**

Returns activity reports for the library.

```
[
  {
    "time": "2023-08-21T12:05:02-0700",
    "duration": 3600,
    "mounts": 0,
    "imports": 0,
    "exports": 0,
    "moves": 0,
    "dataReadByHosts": 0,
    "dataWrittenByHosts": 0,
    "dataWrittenToCartridges": 0,
    "temperatureAverage": 22.0,
    "temperatureMin": 21.0,
    "temperatureMax": 22.5,
    "humidityAverage": 71.0,
    "humidityMin": 71.0,
    "humidityMax": 0.0
  },
  {
    "time": "2023-08-21T11:05:02-0700",
    "duration": 3600,
    "mounts": 0,
    "imports": 0,
    "exports": 0,
    "moves": 0,
    "dataReadByHosts": 0,
    "dataWrittenByHosts": 0,
    "dataWrittenToCartridges": 0,
    "temperatureAverage": 22.0,
    "temperatureMin": 21.0,
    "temperatureMax": 23.0,
    "humidityAverage": 71.0,
    "humidityMin": 69.0,
    "humidityMax": 72.0
  },
  {
    "time": "2023-08-21T10:05:02-0700",
    "duration": 3600,
    "mounts": 0,
    "imports": 0,
    "exports": 0,
    "moves": 0,
    "dataReadByHosts": 0,
    "dataWrittenByHosts": 0,
    "dataWrittenToCartridges": 0,
    "temperatureAverage": 22.0,
    "temperatureMin": 21.0,
    "temperatureMax": 23.0,
    "humidityAverage": 71.6,
    "humidityMin": 69.0,
    "humidityMax": 73.0
  },
  {
    "time": "2023-08-21T09:05:02-0700",
```

```

    "duration": 3600,
    "mounts": 0,
    "imports": 0,
    "exports": 0,
    "moves": 0,
    "dataReadByHosts": 0,
    "dataWrittenByHosts": 0,
    "dataWrittenToCartridges": 0,
    "temperatureAverage": 22.0,
    "temperatureMin": 21.3,
    "temperatureMax": 23.0,
    "humidityAverage": 71.0,
    "humidityMin": 69.0,
    "humidityMax": 72.1
  },
  ...
]

```

Table 6-24 GET /v1/reports/library details

Attributes	Details
time (string)	The time at which the data was recorded.
duration (number)	The number of seconds over which this entry was recorded.
mounts (number)	The number of cartridges mounted to all drives over the duration.
imports (number)	The number of cartridges imported over the duration.
exports (number)	The number of cartridges exported over the duration.
moves (number)	The number of cartridges moved from one location to another over the duration.
dataReadByHosts (number)	The amount of data (in MB) that was read from cartridges by all drives over the duration.
dataWrittenByHosts (number)	The amount of data (in MB) that was written to cartridges from the host by all drives over the duration.
dataWrittenToCartridges (number)	The amount of data (in MB) that was written to cartridges by all drives over the duration. The average compression ratio can be found by dividing <i>dataWrittenByHosts</i> by <i>dataWrittenToCartridges</i> .
temperatureAverage (number)	The average temperature all drives have recorded over the duration (in Celsius), accurate to one decimal place.
temperatureMin (number)	The minimum temperature all drives have recorded over the duration (in Celsius), accurate to one decimal place.

Attributes	Details
temperatureMax (number)	The maximum temperature all drives have recorded over the duration (in Celsius), accurate to one decimal place.
humidityAverage (number)	The average humidity percentage that all drives have recorded over the duration, accurate to one decimal place.
humidityMin (number)	The minimum humidity percentage all drives have recorded over the duration, accurate to one decimal place.
humidityMax (number)	The maximum humidity percentage all drives have recorded over the duration, accurate to one decimal place.

SAS Ports

Use the following commands to view SAS port information:

► **GET /v1/sasPorts**

Returns a list of tape drive SAS ports and their details.

Example 6-33 GET /v1/sasPorts output

```
[
  {
    "location": "sasPort_F1C2R3P0",
    "driveLocation": "drive_F1C2R3",
    "driveSn": "78PG160",
    "portNumber": 0,
    "state": "noDeviceDetected",
    "address": "50050763005a49f6",
    "hashedAddress": 7,
    "speedActual": "1Gbps"
  },
  {
    "location": "sasPort_F1C2R3P1",
    "driveLocation": "drive_F1C2R3",
    "driveSn": "78PG160",
    "portNumber": 1,
    "state": "noDeviceDetected",
    "address": "50050763009a49f6",
    "hashedAddress": 71,
    "speedActual": "1Gbps"
  }
]
```

Table 6-25 GET /v1/sasPorts details

Attributes	Details
location (string)	The location of the SAS port in the library.
driveLocation (string)	The location of the drive the SAS port is attached to.

Attributes	Details
driveSN (string)	The serial number of the drive the SAS port is attached to.
portNumber (number)	The port number on the drive this SAS port represents. This is either 0 or 1.
state (string)	The current state of the SAS port, used to determine host connectivity.
address (string)	The unique identifier for the port.
hashedAddress (string)	A short version of the port address.
speedActual (string)	The speed of data transmission over the port. Possible values are <i>3 Gbps</i> , <i>6 Gbps</i> , and <i>12 Gbps</i> .

► **GET /v1/sasPorts/<location>**

Returns the same attributes as *GET /v1/sasPorts* for a single port.

Slots

Use the following commands to view storage slot information:

► **GET /v1/slots**

Returns a list of storage slots in the library and their details.

Example 6-34 GET /v1/slots output

```
[
  {
    "location": "slot_F1C1R1",
    "state": "normal",
    "contents": [
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null,
      null
    ],
    "puts": 422,
    "putRetries": 6,
    "getRetries": 14,
    "tiers": 9
  },
  {
    "location": "slot_F1C1R2",
    "state": "normal",
    "contents": [
      null,
      null,
```

```

    null,
    null,
    null,
    null,
    null,
    null,
    null,
    null
  ],
  "puts": 257,
  "putRetries": 11,
  "getRetries": 60,
  "tiers": 9
},
{
  "location": "slot_F1C1R3",
  "state": "normal",
  "contents": [
    null,
    null,
    null,
    null,
    null,
    null,
    null,
    null,
    null
  ],
  "puts": 360,
  "putRetries": 0,
  "getRetries": 17,
  "tiers": 9
},
{
  "location": "slot_F1C1R4",
  "state": "normal",
  "contents": [
    "DG I01L5",
    null,
    null,
    null,
    null,
    null,
    null,
    null
  ],
  "puts": 59,
  "putRetries": 0,
  "getRetries": 54,
  "tiers": 9
},
...
]

```

Table 6-26 GET /v1/slots details

Attributes	Details
location (string)	The location of the slot.
state (string)	The current state of the slot. Possible values are <i>normal</i> and <i>inServiceMode</i> .
contents (array of string)	An array of VOLSERS, in order from tier 1 to tier 9. If a tier is empty, its value is <i>null</i> .
puts (number)	The number of times a cartridge has been placed in this slot.
putRetries (number)	The number of times the accessor has had to retry placing a cartridge in this slot, over the lifetime of the slot.
getRetries (number)	The number of times the accessor has had to retry removing a cartridge from this slot, over the lifetime of the slot.
tiers (number)	The number of tiers in this slot.

► **GET /v1/slots/<location>**

Returns the same attributes as *GET /v1/slots* for a single slot in the library

Tasks

Use the following commands to work with libraries and tape drives:

► **GET /v1/tasks**

Returns a list of long-running tasks (active, completed, and failed) for the library, as well as their details.

Example 6-35 GET /v1/tasks output

```
[
  {
    "ID": 366,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "completed",
    "startTime": "2023-08-21T11:07:12-0700",
    "lastUpdateTime": "2023-08-21T11:17:21-0700",
    "percentComplete": 100,
    "duration": 609,
    "user": "",
    "description": "Library Code Update",
    "volser": null
  },
  {
    "ID": 364,
    "type": "inventoryAllTiers",
    "location": "frame_F0",
    "state": "completed",
```

```

    "startTime": "2023-08-18T15:53:42-0700",
    "lastUpdateTime": "2023-08-18T16:06:29-0700",
    "percentComplete": 100,
    "duration": 767,
    "user": "",
    "description": "Inventory with Audit",
    "volser": null
  },
  {
    "ID": 362,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "completed",
    "startTime": "2023-08-18T12:47:50-0700",
    "lastUpdateTime": "2023-08-18T12:58:15-0700",
    "percentComplete": 100,
    "duration": 625,
    "user": "",
    "description": "Library Code Update",
    "volser": null
  },
  {
    "ID": 360,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "completed",
    "startTime": "2023-08-17T13:04:13-0700",
    "lastUpdateTime": "2023-08-17T13:14:18-0700",
    "percentComplete": 100,
    "duration": 605,
    "user": "",
    "description": "Library Code Update",
    "volser": null
  },
  {
    "ID": 358,
    "type": "updateLibraryFirmware",
    "location": "library",
    "state": "failed",
    "startTime": "2023-08-17T11:35:13-0700",
    "lastUpdateTime": "2023-08-17T11:35:24-0700",
    "percentComplete": 0,
    "duration": 11,
    "user": "",
    "description": "Library Code Update",
    "volser": null
  },
  ...
]

```

Table 6-27 GET /v1/tasks details

Attributes	Details
ID (number)	The unique ID of this task.

Attributes	Details
type (string)	A string describing the type of task this entry represents.
location (string)	The location of the hardware that this task affects.
state (string)	The current state of the task.
startTime (string)	The time when the task was initiated.
lastUpdateTime (string)	The last time the task state changed.
percentComplete (number)	A percentage showing how far along the task is. This value is <i>null</i> if the task is not in progress.
duration (number)	The number of seconds the task has been in progress. This value is <i>null</i> if the task is not in progress.
user (string)	The username of whoever initiated the task.
description (string)	A human-readable description of what the task is.
volser (string)	Only used for <i>verifyMedia</i> tasks. This is the VOLSER of the cartridge being verified. For all other tasks, this value is <i>null</i> .

► **GET /v1/tasks/<ID>**

Returns the same attributes as *GET /v1/tasks* for a single long-running task.

► **POST /v1/tasks {"type": "calibrateLibrary"}**

Begins a library calibration task. This request responds with a JSON structure containing the task ID that can be used to query its status. A successful update reports a return code 201.

► **POST /v1/tasks {"type": "inventoryTier0and1", "location": <"library" | "frame_F1">}**

Begins a library inventory that will only scan tiers deeper than tier 1 in the event there is a mismatch from what is expected. Although this has two different options for the location parameter, they are functionally the same. This is because a Diamondback library has only one frame. A successful update reports a return code 201.

► **POST /v1/tasks {"type": "inventoryAllTiers", "location": <"library" | "frame_F1">}**

Begins a full library inventory, making sure to scan every cartridge. Just like the *inventoryTier0and1* task, there are two options for the location parameter. A successful update will report a return code 201.

► **POST /v1/tasks {"type": "startDriveService", "location": <drive location>}**

Places a drive in service mode to begin a service replacement. A successful update reports a return code 201.

► **POST /v1/tasks {"type": "completeDriveService", "location": <drive location>}**

Takes a drive out of service mode and updates drive firmware to the level of the drive that was replaced to complete a service replacement. A successful update reports a return code 201.

- ▶ **POST /v1/tasks {"type": "testDrive", "location": <drive location>}**
Performs a drive test on the specified tape drive. A successful update reports a return code 201.
- ▶ **POST /v1/tasks {"type": "updateDriveFirmware"}**
Update drive firmware at the specified location. The request should contain the *type* JSON structure, followed by an HTTP header with a *Content-Type* of "application/octet-stream" and a *Content-Length* matching the size of the drive firmware image, followed by the binary firmware image itself. A successful update reports a return code 201.

Work Items

These requests are used to move cartridges from one location to another.

- ▶ **POST /v1/workItems {"type": "moveToDrive", "cartridge": <volser>, "sourceInternalAddress": <internal address>, "destinationLocation": <location>, "destinationSN": <serial number>}**
Moves a cartridge to a tape drive. Either the cartridge VOLSER or internal address can be used as the source, and either the drive location or serial number can be used as a destination. This will respond with a JSON structure containing an ID of 0, which will be updated in the future to return a positive value for querying the work item queue. A successful update will report a return code 201.

The progress of this move can be polled by issuing a **GET /v1/dataCartridges/<volser>**, **GET /v1/diagnosticCartridges/<volser>**, or **GET /v1/cleaningCartridges/<volser>** command as appropriate to the cartridge type to monitor its location in the library.

- ▶ **POST /v1/workItems {"type": "moveToIOStation", "cartridge": <volser>, "sourceInternalAddress": <internal address>, "destinationLocation": <location>}**
Moves a cartridge to a slot in the I/O magazine. Either the cartridge VOLSER or internal address can be used as the source, and the I/O slot must be specified as the destination. This will respond with a JSON structure containing an ID of 0, which will be updated in the future to return a positive value for querying the work item queue. A successful update will report a return code 201.

- ▶ **POST /v1/workItems {"type": "moveToSlot", "cartridge": <volser>, "sourceInternalAddress": <internal address>, "destinationLocation": <location>}**
Moves a cartridge to a storage slot. Either the cartridge VOLSER or internal address can be used as the source, and the location of the slot must be specified as the destination. This will respond with a JSON structure containing an ID of 0, which will be updated in the future to return a positive value for querying the work item queue. A successful update reports a return code 201.

The progress of this move can be polled by issuing a **GET /v1/dataCartridges/<volser>**, **GET /v1/diagnosticCartridges/<volser>**, or **GET /v1/cleaningCartridges/<volser>** command as appropriate to the cartridge type to monitor its location in the library.

6.1.3 Parameters

Some commands include optional parameters. There are three parameter classifications:

1. Path parameters
2. Request body parameters
3. Query parameters

Path parameters are added with a slash after the endpoint name and point to a specific resource within that endpoint (**GET /v1/drives/<location>**).

Request body parameters are sent as JSON data and can specify multiple parameters in the request (**POST /v1/tasks** {"type": "inventoryAllTiers", "location": "library"}).

Query parameters are added after a question mark and allow a user to request a subset of the data in a response (**GET /v1/events?after=<timestamp>**).

6.2 Using sg3_utils (Linux)

The sg3_utils Linux package includes the commands necessary to send REST over SCSI request to the library and receive a response to them. The **sg_write_buffer** command is used for sending a request, and the **sg_read_buffer** command is used for receiving a response to it.

The following parameters are needed for the *sg_write_buffer* command to send the request:

Buffer ID: 0x10

Mode: 0x2

Offset: 0 to begin, but it must be incremented by successive command for transferring more than 1 MB of data.

Use the **sg_write_buffer** command to echo a string to the write buffer as a REST request.

```
echo "<string>" | sg_write_buffer --id=0x10 -in=- --mode=2 --offset=0 /dev/<SCSI medium changer device>
```

The following parameters are needed for the **sg_read_buffer** command:

Buffer ID: 0x11

Mode: 0x2 or 0x0

Offset: 0 to begin, but it must be incremented by 1 for each successive request if the amount of data being received from the library is greater than 1MB.

Typical usage of the *sg_read_buffer* command is in the following example:

```
sg_read_buffer -id=0x11 mode=2 -offset=0 -length=0x100000 -raw /dev/<SCSI medium changer device>
```

6.2.1 Example

Example 6-37 shows a *GET /v1/library* example that uses REST over SCSI with *sg3_utils*.

Example 6-36 GET /v1/library command

```
GET /V1/library" | sg_write_buffer --id=0x10 --in=- --mode=2 --offset=0 /dev/sg5 --sg_read_buffer --id=0x11 --mode=2 --length=0x100000 --offset=0 --raw /dev/sg5
```

Example 6-37 Get /v1/library command raw output

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: 767
```

```
[{"name": "ravioli", "status": "driveDegraded", "cartridgeAccess": "normal",
"totalCapacity": 1584, "licensedCapacity": 800, "totalCartridges": 68,
"assignedCartridges": 56, "firmware": "2.9.0.3-11K", "sn": "13B0017", "time":
"2023-08-21T12:31:28-0700", "timezone": "America/Phoenix", "ntpMode": "disabled",
"primaryNtpAddress": null, "secondaryNtpAddress": null, "location": null,
"address": null, "city": null, "state": null, "country": null, "contact": null,
"telephone": null,
"secondaryTelephone": null, "secureCommunications": "enabled",
"autoEjectCleaningCartridges": "enabled", "elasticCapacity": "maxCapacity",
"activeAccessors": "accessorAOnly", "vioStatus": "enabled", "nmaDetection":
"enabled", "capacityUtilThresh": 98.0, "dualAccessorUtilThresh": 98.0}]
```

6.2.2 Endpoints

All endpoints mentioned in this chapter are supported with the `sg3_utils` package.

6.3 Using ITDT for RoS

This section describes how to use IBM Tape Diagnostic Tool (ITDT) for Restover SCSI (RoS) and provides command examples. Using ITDT is the preferred method of using the REST API.

For more information about installing ITDT, see [IBM Tape Diagnostic Tool - Standard Edition - IBM Documentation](#).

ITDT-SE version 9.6.0 or later is required for your operating system. The ITDT package can be from [IBM Support: Fix Central - Select fixes](#).

Run the ITDT commands directly from the command line and include the path to the ITDT folder if you run the command from outside of the ITDT folder.

Command format

Example 6-38 RoS command example

```
ros | rosraw GET | PATCH | POST <url-endpoint> [options] [-o <filename>]
```

Parameters:

- ▶ Use `ros` for formatted output
- ▶ `rosraw` for raw mode, which means no indents, line breaks, or Unicode filtering
- ▶ HTTP-Method:
 - GET
 - PATCH
 - POST
- ▶ `<url-endpoint>` is the RoS method `'/v1/...'`
- ▶ Options in brackets `{}` - string must be masked with single quotation marks (`'`). See Example 6-39 on page 107.
- ▶ Optional file name where the JSON is stored.

The ITDT commands feature the format that is shown in the following examples:


```
./itdt -f /dev/smc30 ros POST /v1/tasks '{"type":"inventoryTier0and1","location":"library"}'  
./itdt -f /dev/smc30 ros GET /v1/logs/LOG_AA004_20230525.zip/export  
./itdt -f /dev/smc30 ros GET /v1/diagnosticCartridges/DG 0200L4'
```

Note: Note: The **ros** parameter in the command string can be replaced with **rosraw** if raw, unformatted data is required.

For more information about using the ITDT command line see, [IBM Tape Diagnostic Tool - Standard Edition](#)

6.3.1 Examples

Example 6-40 shows the **ITDT** command with **ros** to get the library information with formatted output.

Example 6-40 *ros GET /v1/library command*

```
# ./itdt -f /dev/sg3 ros GET /v1/library  
[  
  {  
    "name": "SLT1",  
    "status": "online",  
    "cartridgeAccess": "normal",  
    "totalCapacity": 1232,  
    "licensedCapacity": 800,  
    "totalCartridges": 50,  
    "assignedCartridges": 25,  
    "firmware": "2.10.0.0-029",  
    "sn": "13B0021",  
    "time": "2023-05-25T16:08:40-0500",  
    "timezone": "America/Mexico_City",  
    "ntpMode": "disabled",  
    "primaryNtpAddress": null,  
    "secondaryNtpAddress": null,  
    "location": null,  
    "address": null,  
    "city": null,  
    "state": null,  
    "country": null,  
    "contact": null,  
    "telephone": null,  
    "secondaryTelephone": null,  
    "secureCommunications": "enabled",  
    "autoEjectCleaningCartridges": "enabled",  
    "elasticCapacity": "maxCapacity",  
    "activeAccessors": "accessorA0only",  
    "vioStatus": "enabled",  
    "nmaDetection": "enabled",  
    "capacityUtilThresh": 99.0,  
    "dualAccessorUtilThresh": 98.0
```

```
}  
]
```

Example 6-41 shows the **ITDT** command with **rosraw** to get the library information with unformatted (raw) output.

Example 6-41 rosraw GET /v1/library command

```
# ./itdt -f /dev/sg3 rosraw GET /v1/library  
[{"name": "SLT1", "status": "online", "cartridgeAccess": "normal",  
"totalCapacity": 1232, "licensedCapacity": 800, "totalCartridges": 50,  
"assignedCartridges": 25, "firmware": "2.10.0.0-029", "sn": "13B0021", "time":  
"2023-05-25T16:19:50-0500", "timezone": "America/Mexico_City", "ntpMode":  
"disabled", "primaryNtpAddress": null, "secondaryNtpAddress": null, "location":  
null, "address": null, "city": null, "state": null, "country": null, "contact":  
null, "telephone": null, "secondaryTelephone": null, "secureCommunications":  
"enabled", "autoEjectCleaningCartridges": "enabled", "elasticCapacity":  
"maxCapacity", "activeAccessors": "accessorAOnly", "vioStatus": "enabled",  
"nmaDetection": "enabled", "capacityUtilThresh": 99.0, "dualAccessorUtilThresh":  
98.0}]
```

6.3.2 Endpoints and ITDT

All endpoints mentioned in this chapter are supported with ITDT.

6.4 REST over Ethernet

REST over Ethernet contains all the same endpoints as REST over SCSI. The only difference in accessing them is that a URL must be used with REST over Ethernet, and a user must be authenticated.

A typical RoE URL looks like `https://<IP address>/web/api/v1/<endpoint name>`.

Any existing library user can access the REST API over Ethernet.

To access the API, the endpoints listed in section 6.1.2 can be reached by using the method (GET, PUT, POST) and accessing the URL with the endpoint specified at the end. For example, `GET /v1/library` would be a GET request sent to `https://<IP address>/web/api/v1/library`.

6.4.1 Authentication

The major difference between REST over Ethernet and REST over SCSI (apart from the interface) is that RoE requires a user to be authenticated to issue commands to the library. This uses the same users as the Diamondback web GUI.

The following JSON structure must be sent to the library at `https://<IP address>/web/api/v1/login`: `{"user": <username>, "password": <password>}`

To logout of the session, send the same JSON structure needs to *https://<IP address>/web/api/v1/logout*.

After a user is authenticated, they can continue sending commands to the library over Ethernet until they logout or the session times out due to inactivity based on the library's settings.

6.4.2 Examples

Example 6-42 shows a simple **GET /v1/library** example. Note that rather than list the header attributes like REST over SCSI, REST over Ethernet uses the actual HTTP header to store the error code, content type, and content length.

Example 6-42 GET /v1/library output example

```
URL = 'https://9.11.44.28/web/api/v1/library'
<Response [200]>
[
  {
    "name": "ravioli",
    "status": "driveDegraded",
    "cartridgeAccess": "normal",
    "totalCapacity": 1584,
    "licensedCapacity": 800,
    "totalCartridges": 68,
    "assignedCartridges": 56,
    "firmware": "2.9.0.3-11K",
    "sn": "13B0017",
    "time": "2023-08-21T12:34:52-0700",
    "timezone": "America/Phoenix",
    "ntpMode": "disabled",
    "primaryNtpAddress": null,
    "secondaryNtpAddress": null,
    "location": null,
    "address": null,
    "city": null,
    "state": null,
    "country": null,
    "contact": null,
    "telephone": null,
    "secondaryTelephone": null,
    "secureCommunications": "enabled",
    "autoEjectCleaningCartridges": "enabled",
    "elasticCapacity": "maxCapacity",
    "activeAccessors": "accessorAOnly",
    "vioStatus": "enabled",
    "nmaDetection": "enabled",
    "capacityUtilThresh": 98.0,
    "dualAccessorUtilThresh": 98.0
  }
]
```

6.4.3 Endpoints

All endpoints in this chapter are supported in REST over Ethernet.

6.5 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 indicates a possible user error.
- ▶ 5xx: Server Error. The server takes responsibility for these error status codes.

Examples

On a successful GET request, the library returns an error code of 200, indicating that the request was successful.

Example 6-43 Successful GET request example

```
[
  {
    "name": "ravioli",
    "status": "driveDegraded",
    "cartridgeAccess": "normal",
    "totalCapacity": 1584,
    "licensedCapacity": 800,
    "totalCartridges": 68,
    "assignedCartridges": 56,
    "firmware": "2.9.0.3-11K",
    "sn": "13B0017",
    "time": "2023-08-21T12:34:52-0700",
    "timezone": "America/Phoenix",
    "ntpMode": "disabled",
    "primaryNtpAddress": null,
    "secondaryNtpAddress": null,
    "location": null,
    "address": null,
    "city": null,
    "state": null,
    "country": null,
    "contact": null,
    "telephone": null,
    "secondaryTelephone": null,
    "secureCommunications": "enabled",
    "autoEjectCleaningCartridges": "enabled",
    "elasticCapacity": "maxCapacity",
    "activeAccessors": "accessorAOnly",
    "vioStatus": "enabled",
    "nmaDetection": "enabled",
```

```
    "capacityUtilThresh": 98.0,  
    "dualAccessorUtilThresh": 98.0  
  }  
]
```

On a successful POST request, the library returns a code of 201, which indicates that a task was created.

Example 6-44 Successful post command return code

```
HTTP/1.1 201 Created  
Content-Type: application/json  
Content-Length: 0  
{ "ID": "105" }
```

On a failed request, the library gives some description of why the request failed, such as a 404 error.

Example 6-45 Failed POST request example

```
URL = 'https://9.11.44.28/web/api/v1/librrary'  
<Response [404]>  
{  
  "error_code": 404,  
  "error_msg": "Not Found",  
  "error_type": "Resource not found",  
  "error_description": "URL /web/api/v1/librrary not found"  
}
```



Diamondback self-support

Diamondback is designed to be serviced by the customer with a few customer replaceable units (CRUs). You can view online videos that show the simple repair procedures. Most parts, including the cartridge accessor assembly, can be replaced in 2 minutes or less. Diamondback is also available with full IBM Service, with both Basic and Premium levels of service available.

Self-support example scenarios are described in this chapter including the process and steps to resolve them. These scenarios and solutions can include links to IBM Documentation and the use of service videos available to address the most commonly seen issues and replaced components.

This chapter includes the following topics:

- ▶ 7.1, “Diamondback self-support and serviceability” on page 122
- ▶ 7.2, “Self-support scenario and issue resolution example” on page 125

7.1 Diamondback self-support and serviceability

7.1.1 Self-support and CRUs

Customer Replaceable Units and Field Replaceable Units

The customer is expected to replace parts that are designated as CRUs without the assistance of onsite IBM Service personnel. The field replaceable unit (FRU) parts are replaced by an IBM Service Support Representative (SSR) or an IBM Authorized Warranty Service Provider.

Customer Replaceable Unit

IBM Service provides a CRU to you for you to install if a component must be replaced.

CRU information and replacement instructions can be found on the front door and back door of the Diamondback library. See also 7.1.2, “QR codes” on page 123.

You can use the CRU process to repair your Diamondback library. In situations where a library component failure can be attributed to a CRU, IBM Service sends you the new part. Without the need for any special tools or skills, you can replace the defective part for the new one. All the fasteners from all CRUs in the Diamondback library are captive, which means you do not have to loosen screws. The expected time to replace a CRU in the Diamondback library is less than 2 minutes.

CRU is a simple process with the advantage that the Diamondback library is running in the fastest possible time after the failure.

In the materials that are shipped with a replacement CRU, IBM specifies the return instructions and provides a return label.

Note: You can be charged for the replacement CRU if IBM does not receive the defective CRU within 15 days of your receipt of the replacement.

The following parts in the Diamondback library are CRUs:

- ▶ Robotic accessor
- ▶ Electronics Tray (does not include Compact Flash Card)
- ▶ Compact Flash Card
- ▶ 1600 W Power Supply
- ▶ Drive cannisters
- ▶ Power Distribution Unit (PDU)
- ▶ LTO Service Magazine
- ▶ Y Flex Cable

7.1.2 QR codes

Each of the CRU's for the library has a QR code printed on the frame of the library its self. The front of the library contains the QR codes shown in Figure 7-1.




Figure 7-1 QR codes for items accessed through the front of the Diamondback library

The rear of the Diamondback library frame contains the following QR codes, as shown in Figure 7-2 on page 124, for CRUs that are accessed through the rear door of the library and for installation or movement of the library and for library diagnostics.




Figure 7-2 QR codes for items that are accessed through the rear of the Diamondback library

Each CRU video is also available at the following IBM websites:

- ▶ <https://mediacenter.ibm.com/tag?tagid=diamondback>
- ▶ <https://wwwstage.ibm.com/docs/en/diamondback/2.9.2>

Figure 7-3 shows an example of the ‘how to’ video for Diamondback drive CRU replacement located within the IBM Diamondback Tape Library guide on IBM Docs. Select the Removing Drives option and then click on the play button and the video will start.

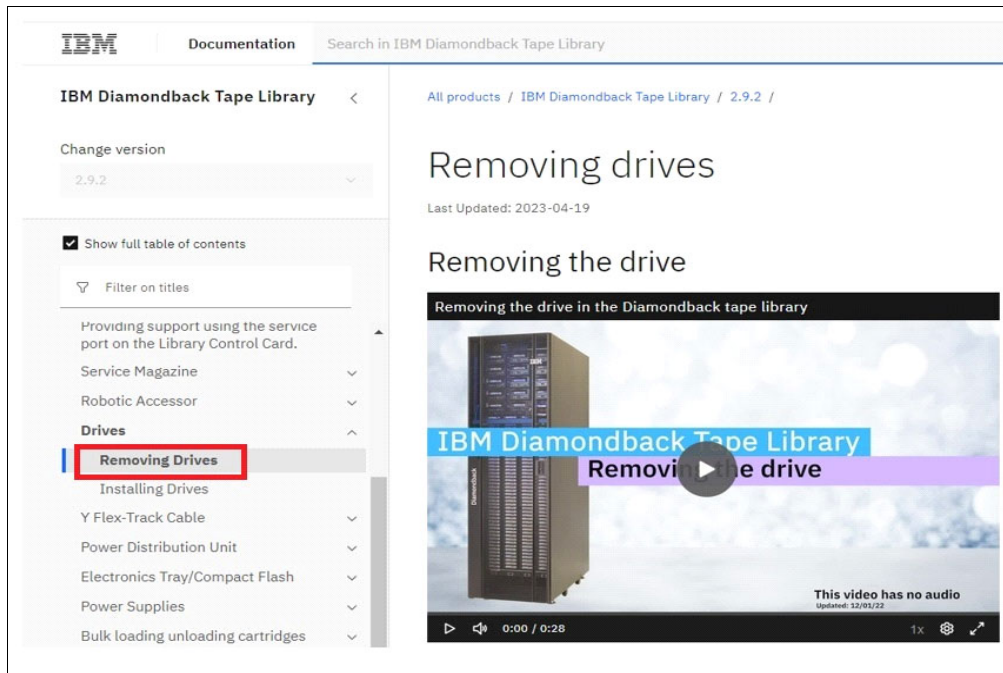


Figure 7-3 IBM docs Diamondback guide

Field replaceable units

When a part that is defective is flagged as a FRU an IBM SSR or an IBM Authorized Warranty Service Provider replaces the part. The replacement part can be directly sent to the customer, or the IBM SSR brings the part to the customer. The SSR contacts the customer to make an appointment to replace the defective part and to arrange, if necessary, access to the data center where the Diamondback library is located.

The following parts in the Diamondback library are FRUs:

- ▶ Backplane docking (BPD) Assembly
- ▶ Front Side Connections (FSC) FRU KIT (includes FSC + BPD-FSC Cable)
- ▶ Canister Docking Interface (CDI) FRU Kit (includes CDI + BPD-CDI Cable)
- ▶ Door switch
- ▶ Cable, front door switch
- ▶ Cable, rear door switch
- ▶ Power switch/cable
- ▶ Left/right light strip FRU kit
- ▶ Y rail
- ▶ Left/right robotic accessor wheel kit
- ▶ LTO Ultra High Density (UHD) Latch assembly
- ▶ LTO UHD slot
- ▶ Front door
- ▶ Rear door
- ▶ Left side cover
- ▶ Right side cover

7.2 Self-support scenario and issue resolution example

This example describes the steps to resolve an issue command-processing issue.

Host cannot process commands.

While running normal backup/restore activities there is a problem reported by the backup application. Figure 7-4 shows an example of using ITDT to manually move a volume and confirm an I/O error is reported when doing a move in the library.

```
C:\ITDT>itdt.exe -f H4-B0-T5-L1 move 260 1025
Moving medium: source 260, destination 1025...
Operation FAILED with errno 1117 The request could not be performed because of an I/O device error

Error Sense Data, Length 96

   0 1  2 3  4 5  6 7  8 9  A B  C D  E F   0123456789ABCDEF
0000 - 7000 0400 0000 0046 0000 0000 4400 0000 [p.....F....D...]
0010 - 0000 0088 0000 0000 1700 0200 3101 1A00 [..@.....1...]
0020 - 0000 0000 0000 0000 0000 0000 0000 0000 [.....]
0030 - 0000 0000 0000 0000 0000 0000 0000 0000 [.....]
0040 - 0000 0000 0000 0000 0000 0000 0000 0000 [.....]
0050 - 0000 0000 0000 5430 4B33 4635 4B37 5856 [.....T0K3F5K7XV]

Exit with code: 1117
```

Figure 7-4 Using ITDT to move a volume results in a failure

To confirm if this is a problem with the library, run the following steps:

Review the [Troubleshooting guide](#) for the [IBM Diamondback Tape Library](#). Per the troubleshooting guide, one of the options is to review if there are events reported in the Management GUI (Figure 5-25 on page 71). See events details in 5.2.3, “Events” on page 69.

Events for errors will have suggestions to resolve the issue as shown in Figure 5-26 on page 71.

For this example, the Diamondback GUI Events table (Figure 7-5) reports an Error event along with another Warning event and Informational events reported at about the same point in time. Because all of them are reported at about the same time, the three of them might be related. Begin by examining the Error event because this has the highest importance.

Location	Type	Time	User	Description	St...	Error Co
✖ Accessor A, Gripper 1	Library	6/12/23, 3:50 PM	Service	Gripper 1 sensor was blocked when it should not have been.		4181
⚠ Accessor A	Library	6/12/23, 3:50 PM	Service	Accessor A detected an obstacle in the accessor or gripper's path		48A0
ℹ Accessor A	Library	6/12/23, 3:50 PM	System	The state of accessor A changed from onlineActive to noMovementAllowed.		0702
ℹ	Library	6/12/23, 3:50 PM	System	The state of the library changed from online to accessorsUnavailable.		0700

Figure 7-5 Diamondback GUI events table

The first steps of the proposed solution in the troubleshooting guide, says to look up the Error event code in the [Events/Fix Procedures](#).

In this case, review the Fix Procedure for Event 4181.

First, the event code and its general description, followed by a detailed description (Figure 7-6). This detailed description provides the possible causes for the problem reported.

4181 - Gripper sensor was blocked when it should not have been

Last Updated: 2023-05-10

The cartridge present sensor in a gripper on the robotic accessor was blocked when it should not have been. This could be caused by a cartridge stuck between the gripper and a slot, drive or service magazine, a faulty sensor on the gripper, or a improperly plugged cable within the robotic accessor. This fix procedure describes how to diagnose and correct the problem and, if needed, remove a stuck cartridge between the robotic accessor and a slot.

Figure 7-6 Gripper error details example

Next, the actual fix procedure begins (Figure 7-7). Step 1 of the procedure says to put the library in Service Mode (Figure 7-8 on page 127). After this is done, open the front door and make a visual inspection for any obstructions and correct them if necessary.

Fix procedure

1. Put the library in service mode:
 - a. From the management GUI, go to **Actions > Change Library State**.
 - b. Choose **In Service Mode (Offline)** option and then click the **Set** button.
2. Open the front door.
3. Clear away any obstacles in the way of the robotic accessor such as a cartridge sticking out of a cell. If a cartridge is sticking out of a cell push it back in.

Figure 7-7 GUI Gripper error fix procedure example

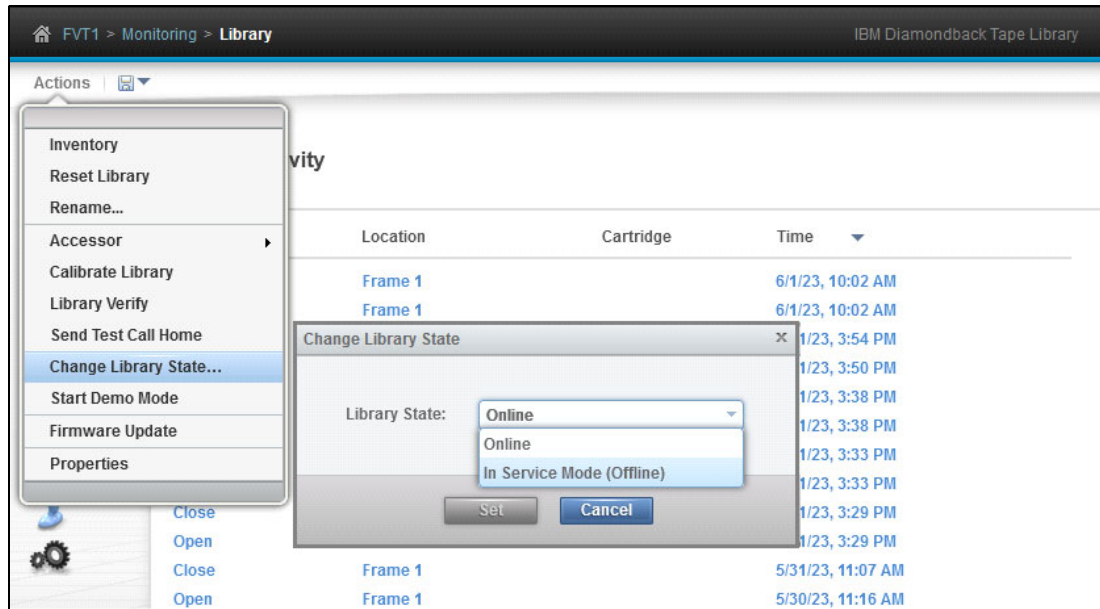


Figure 7-8 Change library status

Step 3 also includes an Additional Troubleshooting section in blue (Figure 7-9). This section contains more detailed and advanced steps that require manual manipulation of the robotic accessor, grippers, or cables. If the administrator is not comfortable with the steps listed, contact IBM and request assistance from IBM support for assistance resolving the problem.




Figure 7-9 Additional troubleshooting steps

If the problem was not resolved by the procedure in step3, the robotic accessor needs to be replaced as described in step4.a (Figure 7-11 on page 128). The replacement procedure steps can be found by scanning the process QR code (Figure 7-10 on page 128) or by following the remove and install videos for the [Robotic Accessor](#).




Figure 7-10 robotic accessor replacement QR code

After the robotic accessor is replaced, continue with steps 4.b through 4.e to bring it online and confirm it works correctly.

4. Replace the robotic accessor:
- Replace the robotic accessor using the [Robotic Accessor](#) videos.
 - Reset the LCC. From the management GUI, go to **Library > Nodes Cards**. Right click on the LCC node card and choose **Reset**. This process will log you out of the GUI.
 - Close the front door and calibrate the library. From the management GUI, go to **Actions > Calibrate Library** and click **Yes**.
 - Put the library online. From the management GUI, go to **Actions > Change Library State** and choose **Online** option.
 - Run Library Verify. From the management GUI, go to **Actions > Library Verify** keep the default options and click **Next**

Figure 7-11 Post robotic accessor replacement steps cont.

For step 5, if the issue was not resolved by fixing the issue or replacing the robotic accessor, contact [IBM support](#) who can assist in further problem determination and resolution.

As noted in the Attention bullet, be sure to monitor the system to ensure the same problem does not reappear. This might mean a different problem is the root cause of the original issue reported. If issues continue to be reported, contact [IBM support](#).

5. If the problem was solved you are done. Otherwise, [contact IBM support](#).

Attention: If the problem has occurred multiple times and still not resolved, [contact IBM support](#).


Figure 7-12 Post robotic accessor replacement steps, resolved

Note: Generally for the self-support procedures, if a CRU is found damaged, replace the CRU by using the [Customer Service Procedures](#).

Damaged Components

During the procedure, if any part is found physically damaged, replace the applicable Customer Replacement Unit (CRU). See [Customer service procedures](#).

Figure 7-13 Damaged components note



Feature codes for IBM Diamondback

The table in this appendix lists the feature codes that are available for the Diamondback Tape Library.

FC	MTM	CSU	Description
0983	4882L9A 4883L9A 4884L9A	N/A	TAA Compliance This feature indicates that the product is TAA compliant.
AM01	4882L9A 4883L9A 4884L9A	N/A	Pre-load media Half Capacity This feature instructs manufacturing to preload media (up to 800 data cartridges as the standard capacity).
AM02	4882L9A 4883L9A 4884L9A	N/A	Pre-load media Full Capacity This feature instructs manufacturing to preload media (up to 1548 data, cleaning, diagnostics cartridges). This feature requires the purchase of a full capacity license.
AM09	4882L9A 4883L9A 4884L9A	N/A	Unified ship LTO media and library order Administrative feature to indicate unified shipment of an LTO9 Media order with an LTO Library frame order. This optional feature instructs the plant to integrate LTO9 Media into select Tape Library frames according to licensed capacity.
1852	4882L9A 4883L9A 4884L9A	No	1X PDU in Top Rack Provides one 1U PDU with 3-phase (wye) or single-phase input power to be mounted in top rack (max of two can be ordered). Each PDU provides six C13 outlets and three C19 outlets.

1853	4882L9A 4883L9A 4884L9A	No	1X PDU in Frame Provides one 1U PDU with 3-phase (wye) or single-phase input power, to be mounted in the Diamondback (max of two can be ordered). Each PDU provides six C13 outlets and three C19 outlets.
1854	4882L9A 4883L9A 4884L9A	No	Customer provided PDU This feature notifies the plant that a customer-provided PDU will be used.
ALH0	4882L9A	No	Expert Care Indicator This feature indicates that a priced Expert Care Service and Support offering is included on the order.
ALJ1	4882L9A	No	1-year Expert Care Basic This feature indicates that a priced Expert Care Basic Service and Support offering is included on the order.
ALJ2	4882L9A	No	2-year Expert Care Basic This feature indicates that a priced Expert Care Basic Service and Support offering is included on the order.
ALJ3	4882L9A	No	3-year Expert Care Basic This feature indicates that a priced Expert Care Basic Service and Support offering is included on the order.
ALJ4	4882L9A	No	4-year Expert Care Basic This feature indicates that a priced Expert Care Basic Service and Support offering is included on the order.
ALJ5	4882L9A	No	5-year Expert Care Basic This feature indicates that a priced Expert Care Basic Service and Support offering is included on the order.
ALL1	4882L9A	No	1-year Expert Care Premium This feature indicates that a priced Expert Care Premium Service and Support offering is included on the order.
ALL2	4882L9A	No	2-year Expert Care Premium This feature indicates that a priced Expert Care Premium Service and Support offering is included on the order.
ALL3	4882L9A	No	3-year Expert Care Premium This feature indicates that a priced Expert Care Premium Service and Support offering is included on the order.
ALL4	4882L9A	No	4-year Expert Care Premium This feature indicates that a priced Expert Care Premium Service and Support offering is included on the order.

ALL5	4882L9A	No	<p>5-year Expert Care Premium</p> <p>This feature indicates that a priced Expert Care Premium Service and Support offering is included on the order.</p>
1553	4882L9A 4883L9A 4884L9A	N/A	<p>Olympus Top Rack</p> <p>This feature code provides a 5U of rack space for client-supplied equipment without requiring any additional floor space in conformance with the Open Compute Project Olympus Rack Specification. It provides 5U of 19-inch rack space with 3 sets of EIA rails and a spacing from the front face of the front EIA rail to the rear face of the middle EIA rail being 745 mm 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. This configuration allows the top rack to be bolted on top of a Diamondback Tape Library. It includes a top cover that allows top access cables to be routed into the top rack. This also includes side air blocks that prevent airflow between the EIA rails and the side covers per the Open Compute Olympus Rack Specification.</p>
1554	4882L9A 4883L9A 4884L9A	N/A	<p>Standard 19 inch Top Rack</p> <p>This feature code provides a 5U of rack space for client-supplied equipment without requiring any additional floor space in conformance with industry standard 19-inch racks. This provides 5U of 19-inch rack space with 2 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. This configuration allows the top rack to be bolted on top of a Diamondback Tape Library. It includes a top cover that allows top access cables to be routed into the top rack.</p>
1555	4882L9A 4883L9A 4884L9A	N/A	<p>Top Rack Front Door</p> <p>This feature provides the front door for the 5U top rack.</p>
1556	4882L9A 4883L9A 4884L9A	N/A	<p>Top Rack Rear Door</p> <p>This feature provides the rear door for the 5U top rack.</p>
AHY2	4882L9A 4883L9A 4884L9A	No	<p>Onsite Code Load</p> <p>This chargeable feature is an indicator that the customer requires someone to physically go onsite to do a code load upgrade. IBM Storage Expert Care Premium includes 2 remote code loads per year. This feature entitles the customer to request them to be performed onsite. The feature quantity must match the number of years Expert Care Premium was ordered for.</p>

AKCH	4882L9A 4883L9A 4884L9A	No	Standard Shipping and Handling This feature indicates that the machine order will be shipped under the 'Standard' Shipping & Handling Service Level Agreement.
AKNC	4882L9A 4883L9A 4884L9A	No	Premium Shipping and Handling This feature indicates that the machine order will be shipped under the 'Premium' Shipping & Handling Service Level Agreement.
1909	4882L9A 4883L9A 4884L9A	Yes	Single Power Source Bifurcated Cable This feature provides a single power outlet connection while maintaining drive redundant power to a Diamondback frame. This feature must be ordered with normal power cord options.
AGK1	4882L9A 4883L9A 4884L9A	Yes	10 m OM3 fiber Cable (LC) This feature provides a 10-meter 50.0/125 micrometer fiber optic cable with LC Duplex connectors.
AGK2	4882L9A 4883L9A 4884L9A	Yes	25 m OM3 fiber Cable (LC) This feature provides a 25-meter 50.0/125 micrometer fiber optic cable with LC Duplex connectors.
AGK3	4882L9A 4883L9A 4884L9A	Yes	80 m OM3 fiber Cable (LC) This feature provides an 80-meter 50.0/125 micrometer fiber optic cable with LC Duplex connectors.
AGKA	4882L9A 4883L9A 4884L9A	Yes	3 m Mini-SAS HD/Mini-SAS HD 4x Cable 3-meter Mini-SAS HD to Mini-SAS HD 4X cable for attachment of up to four Tape Drives (from HBA with SFF-8644 to drives with SFF-8644).
9700	4882L9A 4883L9A 4884L9A	Yes	No Host/SAN Cable from Plant This feature must be placed on any Diamondback Tape Library shipping from the plant with no SAN attachment cables specified when a tape drive is ordered on the Frame. Any required cables must be supplied by the customer.
2704	4882L9A 4883L9A 4884L9A	No	Rackmount Expansion Ethernet Switch for TSSC This feature provides a 24-Port Ethernet switch and an attachment cable for connection to a TSSC for Service. Up to 24 additional connections are provided by this feature.

2715	4882L9A 4883L9A 4884L9A	No	TSSC attachment cable This feature provides a cable to attach a Diamondback to the Ethernet switch provided by the TSSC for Service.
2726	4882L9A 4883L9A 4884L9A	No	Rackmount TSSC This feature provides a 1U rack mounted TSSC console.
1409	4882L9A 4883L9A 4884L9A	Yes	Diagnostic Cartridge for LTO8 and LTO9 This feature provides a L8 CE cartridge. Only one L8 CE cartridge feature #1409 is required in a Diamondback Library.
1411	4882L9A 4883L9A 4884L9A	Yes	Fiber Tape Drive Wrap Tool This feature provides a Fibre wrap tool. Only one Fibre wrap tool feature #1411 is required in a Diamondback Library. The wrap tool is used to isolate drive sled problems.
1412	4882L9A 4883L9A 4884L9A	Yes	SAS Tape Drive Wrap Tool This feature provides a mini SAS HD wrap tool. Only one mini SAS HD wrap tool feature #1412 is required in a Diamondback Library. The wrap tool is used to isolate drive sled problems.
1529	4882L9A 4883L9A 4884L9A	Yes	Preloaded Media enablement kit This feature provides a set of four restraint brackets and accessories for preload media.
ADNR	4882L9A 4883L9A 4884L9A	N/A	No Rules Order Indicator This feature indicates the order will not be validated by e-config. It is used for field install (MES upgrade) orders. Bulk ordering of features for inventory stocking purposes. Note: Only supported for AGH1 and AGH3
AGH1	4882L9A 4883L9A 4884L9A	Yes	LTO 9 FH Fibre Channel Drive This feature is a drive containing one IBM LTO 9 Full-High multimode Fibre Channel Tape Drive with dual port 8 Gbps interface, which can be installed into a Diamondback Tape Library. Cables, ordered separately, are required to attach the Diamondback Tape Library to a host Fibre Channel adapter.

AGH3	4882L9A 4883L9A 4884L9A	Yes	LTO 9 FH SAS Drive This feature is a drive containing one IBM LTO 9 Full-High SAS Tape Drive with a dual port 12 Gbps interface, which can be installed into a Diamondback Tape Library. Cables, ordered separately, are required to attach the Diamondback Tape Library to a host SAS adapter.
8750	4882L9A 4883L9A 4884L9A	Yes	Cleaning Cartridge This feature provides one cleaning cartridge, used to clean the IBM LTO tape drives.
AGL2	4882L9A 4883L9A 4884L9A	Yes	LTO Service Magazine This feature provides an LTO Service Magazine designed for replacing consumed cleaning cartridges or defective data cartridges (maximum 10 cartridges at a time). The service magazine is accessed via the rear door in the same service area as tape drives, cables, power, and electronics.
9948	4882L9A 4883L9A 4884L9A	Yes	3 Phase Power Cord This feature provides a 4.3 m (14 ft) long power cord with an IEC 309 3P+N+G 32A plug, rated for 230 Vac, 24 Amps. This power cord supports 3 phase (wye) power.
9949	4882L9A 4883L9A 4884L9A	Yes	C20-C13 Power Cord for Spare outlets in the Enhanced PDU This feature provides a 2.8 meter long C20-C13 Power Cord to be used between the library power supplies and a 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 customer-supplier PDU. This feature can also be used to power computer equipment from the spare C19 outlets on the Diamondback PDU feature #1852 or #1853.
9954	4882L9A 4883L9A 4884L9A	Yes	Nema L6-30 Power Cord This feature provides a NEMA L6-30 non-watertight 4.3 meter (14 foot) power cord, 200-208, 240 Vac, 24 Amps, used in U.S., Canada, Latin America, and Japan.
9955	4882L9A 4883L9A 4884L9A	Yes	RS 3750DP Power Cord This feature provides a Russellstoll 3750DP Watertight 4.3 meter (14 foot) power cord, 200-208, 240 Vac, 24 Amps, used in U.S, Chicago, Canada, LA, and Japan.

9956	4882L9A 4883L9A 4884L9A	Yes	IEC 309 Power Cord This feature provides a IEC 309, p+n+g, 32A, 4.3 meter (14 foot) power cord, 230 Vac, 24 Amps, used in EMEA.
9957	4882L9A 4883L9A 4884L9A	Yes	PDL 4.3m Power Cord This feature provides a PDL 4.3 meter (14 foot) power cord, 230-240 Vac, 23 Amps, used in Australia and New Zealand
9958	4882L9A 4883L9A 4884L9A	Yes	Korean 4.3m Power Cord This feature provides a non-watertight 4.3 meter (14 foot) power cord, 200-208, 240 Vac, 24 Amps, with Korean plug SJ-3302, used in North and South Korea.
9970	4882L9A 4883L9A 4884L9A	Yes	Dual 4.3M Power Cord This feature provides dual 4.3 meter (14-foot) 250 volt (V) AC 16 Amp power cords with IEC 309 Type 2P+GND 16A connectors, Hubbell type HBL316P6W for connection to Hubbell type HBL316R6W receptacles. This is the default power cord if no other feature is specified, in all countries except Argentina, Australia, Brazil, Canada, China, Japan, Korea, New Zealand, Philippines, South Africa, Taiwan, and the United States. If feature code #1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.
9972	4882L9A 4883L9A 4884L9A	Yes	Dual 4.3M Power Cord NON-WT This feature provides dual 4.3 meter (14-foot) 250 volt (V) AC 15 Amp power cords with a non-watertight connector NEMA L6-15P for connection to NEMA L6-15R receptacles. This is the default power cord if no other feature is specified, in the United States, Canada, and Philippines. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.
9976	4882L9A 4883L9A 4884L9A	Yes	Dual 4.3M Power Cord Argentina This feature provides dual (2x) 4.3 meter (14-foot) 10 Amp 250 Vac power cords with an IRAM 2073 plug for Argentina. This is the default power cord for Argentina if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.

9977	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord Brazil</p> <p>This feature provides dual (2x) 4.3 meter (14-foot) 15 Amp 250 Vac power cords with an Earth Pin InMetro NBR 14136 plug for Brazil. This is the default power cord for Brazil if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9978	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord Australia/New Zealand</p> <p>This feature provides dual (2x) 4.3 meter (14-foot) 10 Amp 250 Vac power cords with an AS/NZS 3112/2000 plug for Australia and New Zealand. This is the default power cord for Australia and New Zealand if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9979	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord Japan</p> <p>This feature provides dual (2x) 4.3 meter (14-foot) 15 Amp 250 Vac power cords with a JIS C8303, C8306 plug for Japan. This is the default power cord for Japan if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9980	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord China</p> <p>This feature provides dual (2x) 4.3 meter (14-foot) 10 Amp 250 Vac power cords with a GB 2099.1, 1002 plug for China. This is the default power cord for China if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9981	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord Korea</p> <p>This feature provides dual (2x) 4.3 meter (14-foot) 12 Amp 250 Vac power cords with an earth pin KS C8305, K60884-1 plug for Korea. This is the default power cord for Korea if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>

9982	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord Taiwan</p> <p>This feature provides dual (2x) 4.3 meter (14-foot) 10 Amp 250 Vac power cords with a CNS 10917-3 plug for Taiwan. This is the default power cord for Taiwan if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9983	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord So. Africa</p> <p>This feature provides dual (2x) 4.3 meter (14-foot) 16 Amp 250 Vac power cords with a SANS 164-1 plug for South Africa. This is the default power cord for South Africa if no other feature is specified. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9984	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord NEMA L6/20P</p> <p>This feature provides dual 4.3 meter (14-foot) 250 volt (V) AC 15 Amp power cords with NEMA L6-20P non-watertight plugs for connection to NEMA L6-20R receptacles. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9985	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord Russellstoll</p> <p>This feature provides dual 4.3 meter (14-foot) 250 volt (V) AC 15 Amp power cords with Russellstoll 3720DPU2 watertight connectors for connection to Russellstoll 3743U2 or 9R23U2W receptacles, or Russellstoll 3913U2 or 9C23U2 connectors. If feature code 1909 (Single Power Source Bifurcated Cable) is also on your order, this feature provides one of the above 4.3 meter power cords.</p>
9989	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 4.3M Power Cord PDU to PSU</p> <p>This feature provides two power cords 4.3 meters (14 feet) in length, rated at 250 Vac, 10 amps with IEC C14 plugs. These power cords can be plugged into the spare outlets on (a) IBM provided PDUs (Feature codes #1851 or #1853) that are installed in an adjacent library or (b) customer-provided PDUs that are mounted external to the library. One or two adjacent libraries can be powered by the spare outlets available on the PDUs. The adjacent library to be powered can be a maximum of three frames away from the library containing the PDUs. By using this feature, up to three libraries can be powered by the two customer facility outlets that feed the PDUs.</p>

9990	4882L9A 4883L9A 4884L9A	Yes	<p>Dual 2.8m power cord PDU to PSU</p> <p>This feature provides two power cords 2.8 meters (9.18 feet) in length, rated at 250 Vac, 10 amps with IEC C14 plugs. For Diamondback this power cord is compatible with feature code #1854.</p>
1600	4882L9A 4883L9A 4884L9A	Yes	<p>Full Capacity License</p> <p>This feature provides a license key to increase storage from half capacity (up to 800 data cartridges) to full capacity. Half Capacity is standard. Max licensed capacity is 1584 slots for LTO. A license key is needed in order to reach full capacity through a key label.</p>
1856	4882L9A 4883L9A 4884L9A	Yes	<p>Overhead Access Ladder</p> <p>This feature provides top rack access. One overhead ladder is recommended per data center.</p>

Related publications and links

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.

Service Video Library

The following link provides access to a number of videos related to deploying, using and servicing the IBM Diamondback library. This is information that goes beyond the QR codes present in this document and on the library its self.

<https://mediacenter.ibm.com/channel/Diamondback+Tape+Library/289311162>

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 Storage Archive Enterprise Edition V1.3.2.2: Installation and Configuration Guide, SG24-8333 or IBM Storage Archive Single Drive Edition SG24-8090
- ▶ IBM Linear Tape File System Enterprise Edition V1.1.1.2: Installation and Configuration Guide SG24-8143
- ▶ 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:

<https://ibm.com/redbooks>

Publications

IBM Diamondback user guide:

<https://www.ibm.com/docs/en/diamondback/2.9.2>

LTO 9 drive information:

<https://www.ibm.com/products/lto-9-tape-drive>

LTO 9 for Hybrid Cloud:

<https://www.ibm.com/downloads/cas/1W2J85JX>

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>

Redbooks

IBM Diamondback Tape Library Guide

SG24-8548-00

ISBN DocISBN



(0.5" spine)

0.475" x 0.873"

250 <-> 459 pages



SG24-8548-00

ISBN 0738459410

Printed in U.S.A.

Get connected

