HITACHI
Inspire the Next

Hitachi Tiered Storage Manager Software
User's Guide

FASTFIND LINKS

Document Organization
Product Version
Getting Help
Table of Contents

Hitachi Data Systems

MK-94HC090-12

Contents

Preface xi

Intended Audience xii
Product Version xii
Release Notes xii
Document Revision Level xii
Document Organization xiii
Referenced Documents xiii
Document Conventions xiv
Convention for Storage Capacity Values xv
Getting Help xvi
Comments xvi

1 Tiered Storage Manager and migration. 1-1

Migration using Tiered Storage Manager 1-2
Purpose of migration 1-2
Meeting SLOs 1-2
Correcting a performance bottleneck 1-3
Responding to an organizational change 1-3
Migration steps 1-3
Planning for migration 1-3
Preparing the migration environment. 1-4
Performing the migration 1-5

2 Using Tiered Storage Manager 2-1

Determining domain controllers 2-2
Managing volumes in tiers 2-2
Using the free space in a storage subsystem to create volumes 2-5
Adding storage subsystems 2-5
Refreshing storage subsystem information 2-6
Connecting external storage subsystems 2-6
Grouping target volumes 2-7

Setting rules for migration 2-7
Array group avoidance rules 2-8
Array group selection rule 2-8
Preparing to migrate volumes 2-10
Volume availability for migration 2-12
Volumes specifiable as migration targets when the migration source volume is a LUSE 2-12
Filtering the number of candidate volumes 2-13
Migrating data to DP pools 2-14
Using pool search conditions to create storage tiers 2-15
Rules for excluding pools 2-15
Discarding zero-data after migrating to DP pools 2-16
Erasing data after migration 2-16
Creating and executing tasks to check the execution status 2-17
Migration tasks 2-19
Locking and unlocking tasks 2-19
Shredding tasks 2-21
Volume creation tasks 2-21
External mapping tasks 2-22
Notes on tasks 2-22
Updating the Tiered Storage Manager database 2-24
Purpose of refreshing 2-24
Automatic refreshing 2-25
Refresh status indicators 2-26
Searching volumes, pools, and free space 2-27
Condition elements of volume filter conditions 2-27
Searching for volumes in a storage domain 2-27
Searching for volumes in a storage subsystem registered in Device Manager 2-29
Condition elements of pool filter conditions 2-29
Condition elements of free space filter conditions 2-29
Notes on filter conditions 2-30
Allocating host paths to volumes 2-30
Using performance information acquired from Tuning Manager 2-31
Connecting to Tuning Manager 2-31
Acquiring array group performance information 2-31
Using mainframe volumes 2-32

3 Using the Web client 3-1

About the Tiered Storage Manager Web client 3-2
Structure of the windows of the Tiered Storage Manager Web client 3-2
Notes on GUI operations 3-3
Login and logout 3-4
Logging in 3-4
Logging out 3-5
Registering licenses 3-5

Setting up user accounts 3-5
User permissions 3-6
Adding a user 3-8
Setting permissions for a user 3-9
Viewing or editing a user profile 3-10
Viewing or editing another user's profile 3-10
Viewing or editing your own profile 3-10
Changing passwords 3-10
Changing another user's password 3-11
Changing your own password 3-11
Setting password conditions 3-12
Setting the automatic locking count for user accounts 3-12
Changing the lock status of a user 3-13
Deleting a user 3-14
Managing user authentication methods 3-15
Linking users to an external authentication server 3-15
Changing the authentication method of a user 3-16
Enabling external authentication using an external authorization server 3-17
Registering an authorization group 3-17
Setting permissions for an authentication group 3-18
Deleting an authorization group 3-18
Setting a message for the warning banner 3-19
Setting up storage domains 3-20
Creating a storage domain 3-20
Checking a storage domain's properties 3-20
Editing storage domain properties 3-20
Deleting a storage domain 3-21
Refreshing storage domains 3-22
Setting up storage tiers 3-22
Creating a storage tier by using the Create Tier dialog box 3-22
Creating a storage tier by specifying volume filter conditions 3-24
Assigning cost attributes to a storage tier 3-25
Checking a storage tier's properties 3-26
Editing a storage tier's properties 3-27
Deleting a storage tier 3-28
Creating volumes 3-28
Searching for free space and creating volumes 3-28
Adding and refreshing storage subsystems 3-29
Registering storage subsystems in Device Manager 3-29
Refreshing storage subsystems 3-29
Connecting external storage subsystems to a storage domain 3-29
Setting up external mapping 3-30
Setting up migration groups 3-31
Creating a target volume group by using the Create MG dialog box 3-31
Creating a volume group by filtering volumes 3-32
Checking a migration group's properties 3-34
Editing a migration group's properties 3-34

Deleting a migration group . . . 3-35
Performing migration . . . 3-36
Migrating volumes by using the Migration Wizard . . . 3-36
Migrating volumes by using the Migrate MG dialog box . . . 3-39
Working with tasks . . . 3-40
Executing tasks . . . 3-40
Checking the execution status of tasks . . . 3-41
Searching for tasks . . . 3-41
Stopping executed tasks . . . 3-41
Canceling tasks . . . 3-42
Deleting tasks . . . 3-42
Editing task properties . . . 3-43

4 Troubleshooting . . . 4-1

Responding to error conditions . . . 4-2
Troubleshooting on the Web client . . . 4-3
Viewing log data . . . 4-9
Message logs . . . 4-9
Trace logs . . . 4-9

A Reference Information . . . A-1

Supported storage subsystems . . . A-2
Storage subsystems that can be registered as domain controllers . . . A-2
External storage subsystems . . . A-4
Storage subsystems in which volumes can be created . . . A-6
Volumes as sources or targets . . . A-6
Volumes whose usability as migration sources or targets can be determined A-6
Volumes whose usability as migration sources or targets cannot be determined . . . A-10
Volume filter conditions . . . A-11
Filter conditions available for existing tiers and domains . . . A-11
Filter conditions for specific subsystems . . . A-15
Input rules for filter conditions . . . A-16
Pool filter conditions . . . A-17
Suggested maximum operational usage . . . A-18

B Details of the array group selection rules . . . B-1

BalanceCapacity . . . B-2
MaximumCoverage . . . B-2
MinimumCoverage . . . B-4

C Relationship of terminology for different storage subsystems . . . C-1

Tiered Storage Manager terminology . . . C-2

D Notes about using earlier versions of the Web client D-1
Searching for tasks by specifying a storage subsystem name D-2
Volumes displayed as candidates for migration targets D-2
Message displayed when storage domain information must be refreshed D-2
The number of candidate volumes to be displayed D-2
Change in the default setting for deleting data D-3
Display of task owner names D-3

Glossary

Acronyms and Abbreviations

Index

Preface

This document describes Hitachi Tiered Storage Manager Software and ways that you can use it to improve management of storage systems.

Notice: The use of Tiered Storage Manager and all other Hitachi Data Systems products is governed by the terms of your agreement(s) with Hitachi Data Systems.

This preface includes the following information:

- Intended Audience
- Product Version
- Release Notes
- Document Revision Level
- Document Organization
- Referenced Documents
- Document Conventions
- Convention for Storage Capacity Values
- Getting Help
- Comments

Intended Audience

This document is written for storage administrators who should have:

- Basic knowledge of storage area networks (SANs)
- Basic knowledge about management tools specific to storage subsystems
- Basic knowledge about the operating system hosting Hitachi Tiered Storage Manager, such as Windows or Solaris
- Knowledge about system operation and management tasks using Hitachi Device Manager (Hitachi Device Manager installation, user settings, user group settings, volume (LDEV) creation, and logical group creation)

Product Version

This document revision applies to Hitachi Tiered Storage Manager version 6.4.

Release Notes

Release notes can be found on the documentation CD. Release notes contain requirements and more recent product information that may not be fully described in this manual. Be sure to review the release notes before installation.

Document Revision Level

Revision	Date	Description
MK-94HC090-P	April 2005	Preliminary Release
MK-94HC090-01	November 2005	Revision 1, supersedes and replaces MK-94HC090-P
MK-94HC090-02	April 2006	Revision 2, supersedes and replaces MK-94HC090-01
MK-94HC090-03	October 2006	Revision 3, supersedes and replaces MK-94HC090-02
MK-94HC090-04	June 2007	Revision 4, supersedes and replaces MK-94HC090-03
MK-94HC090-05	September 2007	Revision 5, supersedes and replaces MK-94HC090-04
MK-94HC090-06	January 2008	Revision 6, supersedes and replaces MK-94HC090-05
MK-94HC090-07	May 2008	Revision 7, supersedes and replaces MK-94HC090-06
MK-94HC090-08	February 2009	Revision 8, supersedes and replaces MK-94HC090-07
MK-94HC090-09	July 2009	Revision 9, supersedes and replaces MK-94HC090-08
MK-94HC090-10	November 2009	Revision 10, supersedes and replaces MK-94HC090-09
MK-94HC090-11	December 2009	Revision 11, supersedes and replaces MK-94HC090-10

Revision	Date	Description
MK-94HC090-12	June 2010	Revision 12, supersedes and replaces MK-94HC090-11

Document Organization

The following table provides an overview of the contents and organization of this document. Click the chapter title in the left column to go to that chapter. The first page of each chapter provides links to the sections in that chapter.

Chapter	Description
Tiered Storage Manager and migration	Describes migration performed using Tiered Storage Manager.
Using Tiered Storage Manager	Describes the knowledge you need to use Tiered Storage Manager to work with volumes.
Using the Web client	Describes how to log in to the Tiered Storage Manager Web client to register users and migrate data on volumes.
Troubleshooting	Explains the causes of errors that occur during the use of the Tiered Storage Manager Web client, and what to do to resolve those errors.
Reference Information	Describes the volumes that cannot be used as migration sources or migration targets in Tiered Storage Manager. It also explains the character strings displayed in the Migration Restrictions field.
Details of the array group selection rules	Describes the order in which Tiered Storage Manager selects migration target volumes, based on the array group selection rule, which is set as part of the LDEV selection rules.
Relationship of terminology for different storage subsystems	Lists the relationship between Hitachi storage subsystem terms and Storage Networking Industry Association (SNIA) terms.
Notes about using earlier versions of the Web client	Explains the differences between using an earlier version of the Web client and using the most recent version.
Glossary	Defines the special terms used in this document.
Acronyms and Abbreviations	Defines the acronyms and abbreviations used in this document.
Index	Lists the topics in this document in alphabetical order.

Referenced Documents

The following Hitachi referenced documents can be found on the applicable Hitachi documentation CD:

- Hitachi Storage Command Suite Documents:
  - Hitachi Tiered Storage Manager Server Configuration and Operation Guide, MK-08HC158
  - Hitachi Tiered Storage Manager CLI Reference Guide, MK-94HC091
  - Hitachi Tiered Storage Manager Messages, MK-94HC092

- Hitachi Device Manager Server Configuration and Operation Guide, MK-08HC157
- Hitachi Storage Command Suite Server Installation Guide Device Manager Provisioning Manager Tiered Storage Manager, MK-98HC150
- Hitachi Device Manager Agent Installation Guide, MK-92HC019
- Hitachi Device Manager Error Codes, MK-92HC016
- Hitachi Device Manager Mainframe Agent User's Guide, MK-96HC130
- Hitachi Tuning Manager Agent Administration Guide, MK-92HC013
- Hitachi Tuning Manager Hardware Agent Reference, MK-95HC111
- Hitachi Tuning Manager Operating System Agent Reference, MK-95HC112
- Hitachi Tuning Manager Server Administration Guide, MK-92HC021

- Hitachi Enterprise Storage Systems Documents:
  - Virtual Partition Manager User's Guide, MK-94RD259
  - Data Retention Utility User's Guide, MK-94RD210
  - Universal Storage Platform™ and Network Storage Controller LUN Expansion, Virtual LVI/LUN and Volume Shredder User's Guide, MK-94RD205
  - Hitachi TagmaStore Universal Storage Platform and Hitachi TagmaStore Network Storage Controller Universal Volume Manager User's Guide, MK-94RD220

Document Conventions

This document uses the following typographic conventions:

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click **OK**.
Italic	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: `copy` *source-file target-file* *Note:* Angled brackets (< >) are also used to indicate variables.
`screen/code`	Indicates text that is displayed on screen or entered by the user. Example: `# pairdisplay -g oradb`
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: `# pairdisplay -g <group>` *Note:* Italic font is also used to indicate variables.
[ ] square brackets	Indicates optional values. Example: [a | b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a | b } indicates that you must choose either a or b.

Convention	Description
| vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a | b] indicates that you can choose a, b, or nothing. { a | b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to important and/or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (e.g., disruptive operations).
	WARNING	Warns the user of severe conditions and/or consequences (e.g., destructive operations).

Convention for Storage Capacity Values

Physical storage capacity values (e.g., disk drive capacity) are calculated based on the following values:

Physical Capacity Unit	Value
1 KB	1,000 bytes
1 MB	$1{,}000^2$ bytes
1 GB	$1{,}000^3$ bytes
1 TB	$1{,}000^4$ bytes
1 PB	$1{,}000^5$ bytes
1 EB	$1{,}000^6$ bytes

Logical storage capacity values (e.g., logical device capacity) are calculated based on the following values:

Logical Capacity Unit	Value
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1{,}024^2$ bytes
1 GB	1,024 MB or $1{,}024^3$ bytes
1 TB	1,024 GB or $1{,}024^4$ bytes
1 PB	1,024 TB or $1{,}024^5$ bytes
1 EB	1,024 TB or $1{,}024^6$ bytes

Logical Capacity Unit	Value
1 BLOCK	512 BYTES

Getting Help

The Hitachi Data Systems Support Center staff is available 24 hours a day, seven days a week. To reach us, please visit the support Web site for current telephone numbers and other contact information:

http://www.hds.com/services/support/. If you purchased this product from an authorized HDS reseller, contact that reseller for support.

Before calling the Hitachi Data Systems Support Center, please provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any error message(s) displayed on the host system(s).

Comments

Please send us your comments on this document: doc.comments@hds.com. Include the document title, number, and revision and refer to specific section(s) and paragraph(s) whenever possible.

Thank you! (All comments become the property of Hitachi Data Systems Corporation.)

1

Tiered Storage Manager and migration

This chapter describes the migration that is performed by Tiered Storage Manager. For detailed information necessary to perform migration, see Using Tiered Storage Manager. For migration procedures, see Using the Web client.

- Migration using Tiered Storage Manager
- Purpose of migration
- Migration steps

Migration using Tiered Storage Manager

Tiered Storage Manager is software that is used to perform migration. The term *migration* refers to moving the data stored on one volume to another volume. Tiered Storage Manager moves the data on a predefined set of volumes to another set of volumes that have the same characteristics.

The set of migration source volumes containing data is called a *migration group*. The set of target volumes to which the data is migrated is called a *storage tier*. Volumes that make up a storage tier, including virtual volumes, must be in the same storage subsystem. This storage area is called a *storage domain*. A storage subsystem that has a storage area defined as a storage domain is called a *domain controller*.

One of the benefits of Tiered Storage Manager is that, although the migration performed by Tiered Storage Manager changes the volumes data is stored on, data access from application servers is not blocked either during or after the migration.

Note: Data still remains on the migration source volumes after a migration. To ensure security, you can use a shredding task to erase data from the migration source volumes. You can also lock volumes so that any data that you are legally obliged to keep cannot be tampered with.

Purpose of migration

As storage data ages, requirements concerning both storage and retrieval of data can change. These requirements can be in the form of performance Service Level Objectives (SLO) or changes in an organizational structure that result in changes that are required in how data is handled and stored—the structure of storage systems.

Following are the *three most common reasons* for migration:

- A change in SLO that either allows for a decrease in storage costs or requires increased performance.
- A performance bottleneck in the subsystem caused by high demands on resources within the same array group, when port and cache problems have been eliminated as a cause.
- Organizational changes that require the data center to non-disruptively restructure its data storage.

Meeting SLOs

You can use Tiered Storage Manager to move data at any given time to volumes that satisfy the SLOs for that data.

While much data must be stored over long periods of time, depending on business needs and legal requirements, many types of data can experience a few rapid changes in need for access and therefore SLOs over the life cycle of the data.

In some cases, access requests to large amounts of incoming data declines after the first few weeks or months (as in email), and in other cases, a specific event might result in a flood of requests for a small set of data (as in audio downloads for a popular entertainer after a large concert).

As email arrives, recipients typically view each message at least once. Later, they will selectively search for and review specific email messages. After several weeks, access demands are more sporadic, and yet the data must continue to be stored. This rapid decrease in email message viewing over time provides an opportunity to reduce the cost of storage.

Email storage for an enterprise typically contains multiple volumes on which the email data of all employees is stored. To handle and distribute the load, volumes storing recent emails can be placed in array groups that have a disk-rotational speed of 12,000 rpm. For incoming email data that arrives in the month of May, requests for viewing vastly diminish by July. So, in July, the volumes containing May data can be migrated, as a set, to volumes in a less expensive array group rotating at 8,000 rpm, and still handle access requests well within the SLO.

Correcting a performance bottleneck

Another reason that migration can become necessary is a performance bottleneck on the subsystem. In some cases, too many people might be attempting to gain access to resources that reside within the same array group. This situation might arise in response to cyclic events like end-of-year accounting needs or one-of-a-kind events related to demands for media in response to entertainment events or disasters.

Responding to an organizational change

As organizations change, the applications they use can either change or be upgraded. These changes are reflected by either changing storage requirements or changes in the storage environment. In these cases, migration might be required either to maintain SLOs or as part of a system integration.

Migration steps

Migration requires three steps: planning, preparation of the environment, and performing the migration.

Planning for migration

To plan for migration:

1. Identify the set of volumes (migration group) that you want to migrate together.

   Because the volumes are migrated at the same time, it is preferable that the set of volumes have the same SLO. In most cases, the data used by a specific application is used.
2. Determine the SLO required for the volumes.

Determine the SLO based on the characteristics of the applications and tasks that use the data. When determining the SLO, consider establishing SLOs for the storage attributes, such as performance and reliability. The following shows an example of the relationship between the storage attributes and SLO perspectives.

SLO perspective	Storage attributes
Performance management	• Rotational speed of the disk • Usage of the target array group (amount of access) • Subsystem type • RAID level
Reliability management	• RAID level • Disk type (such as FC, SATA, BD, or SAS) • Number of years the disk has been used
Reduction of storage operation costs	• Subsystem type (power consumed by the target subsystem) • Status of empty volumes in the storage for the entire system

3. Identify the destination volume or storage tier (migration target) and when to perform migration.

   Migration should be performed when an SLO changes or cannot be maintained. Based on this, you can determine the migration target and when migration should be performed.

When planning has been completed, prepare the migration environment.

Preparing the migration environment

For the detailed knowledge necessary to perform migration, see Tiered Storage Manager and migration. For the migration procedure, see Using the Web client.

1. Define a storage domain.

   In Tiered Storage Manager, register as storage domains the storage subsystems containing the volumes to be used in migration groups and storage tiers.

2. Based on the SLO perspective determined in step 2 of Planning for migration, categorize the existing volumes (create storage tiers). If there are not enough volumes in the storage tier, create additional volumes.

   For example, if the SLO perspective is performance and you want to categorize volumes by disk type (FC or SATA), you need to sort the volumes into FC and SATA groups.

Note: A volume can belong to multiple storage tiers at the same time. Accordingly, the SLO differs according to the application, but you can nevertheless define a storage tier for each aspect.

3. Define migration groups as described in step 1 of Planning for migration

After the migration groups and storage tiers are defined, you can perform migration at the desired time.

Performing the migration

The volume locking, unlocking, and data erasure tasks described in steps 5, 6, and 7 can be performed by using the CLI client.

1. Define the correspondence between the migration source volumes and the migration target volumes (volume pairs).

   This definition is registered as a migration task.

2. Execute the migration task.

   Execute the migration task at the time defined in step 3 of Planning for migration. You can also execute the migration task immediately after you have defined volume pairs in step 1.

   When you create a migration task, if you specify that the source data be erased, data is automatically deleted from the migration source volumes after the data is migrated.

3. In step 2, if you did not specify the source data be erased, erase (shred) the data if necessary.

4. Check the status of the migration task.

   Check the status of the executed migration task, and take appropriate action if a problem occurs.

5. If necessary, lock the volumes.

   You can create and execute a locking task to prevent anyone from tampering with the data stored in the migration group. When creating the locking task, specify how long the volumes are to be locked. Use the CLI to create the locking task.

6. If necessary, unlock the volumes.

   When the lock period you specified in step 5 expires, you can create and execute an unlocking task. Use the CLI to create the unlocking task.

7. If necessary, erase (shred) the data.

   To erase data that does not need to be stored, create and execute a shredding task. Use the CLI if you want to erase data manually, rather than specifying, when you created the migration task, that data be automatically erased.

2

Using Tiered Storage Manager

This chapter provides the information you need to use Tiered Storage Manager to work with volumes. For details about how to use the Web client, see About the Tiered Storage Manager Web client and additional information in Using the Web client.

- Determining domain controllers
- Managing volumes in tiers
- Using the free space in a storage subsystem to create volumes
- Connecting external storage subsystems
- Grouping target volumes
- Preparing to migrate volumes
- Creating and executing tasks to check the execution status
- Updating the Tiered Storage Manager database
- Searching volumes, pools, and free space
- Allocating host paths to volumes
- Using performance information acquired from Tuning Manager
- Using mainframe volumes

Determining domain controllers

When you install Tiered Storage Manager, your first step is to register the storage subsystems to be used as domain controllers in Tiered Storage Manager. The storage area that Tiered Storage Manager can manage by using a domain controller is referred to as a *storage domain*. For detailed information on supported subsystems, see Supported storage subsystems.

Managing volumes in tiers

By specifying filter conditions based on a volume's characteristics, you can group volumes to be used as migration target candidates. Each group is called a *storage tier*. You can add filter conditions until the necessary storage tier has been created. If the number of volumes that match the filter conditions changes after a storage tier has been created, you can automatically reflect the changes in the storage tier by refreshing Tiered Storage Manager. If volumes in a storage tier are deleted from a storage subsystem, refreshing Tiered Storage Manager also deletes the volumes from the storage tier.

For example, you could create the following storage tiers by using the indicated filter conditions:

- Create a storage tier by specifying an externally connected storage subsystem.

  To migrate infrequently used data from a domain controller to an externally connected storage subsystem, specify **Subsystem** or **Subsystem Serial Number** in the filter conditions for creating the storage tier.

- Create a storage tier by specifying an array group.

  When an array group contains multiple volumes that are frequently accessed resulting in degraded response to hosts, you might want to migrate the data into volumes in another array group. In this case, specify **Controller Array Group** or **Array Group** in the filter conditions for creating the storage tier.

- Create a storage tier by specifying disk attributes.

  To migrate data to a faster disk for quicker response to hosts, specify **Disk Type** or **Disk RPM** in the filter conditions for creating the storage tier.

In addition, you can create the following storage tiers based on performance information obtained from Tuning Manager:

- Create storage tiers based on the usage of array groups.

  If the performance information from Tuning Manager indicates that an array group contains high-usage volumes, you can design a migration plan that aims for balanced loads by creating a storage tier that uses volumes in array groups with lower usage rates. To create a storage tier composed of array groups that have different busy rates, specify **Array Group Busy Rate** or **Array Group Max Busy Rate** in the filter conditions.

If you use virtual volumes, you can also create the following storage tier:

- Create a storage tier by specifying a DP pool volume or DP pool.

  When you are using the Dynamic Provisioning feature of Universal Storage Platform V/VM, you can specify a DP pool volume or DP pool as a migration target. When you specify a DP pool as the migration target, a DP pool volume that has the same size and emulation type as the source volume is created in the DP pool when migration is performed. To dynamically create a DP pool volume in the DP pool when performing migration, use either of the following methods to create the storage tier:

  - Specify **Volume(s)** as the filter type, Dynamic Provisioning as the filter condition, and **DP-Pool-VOL** as the value.
  - Specify **Pool** as the filter type and **Free Capacity** of the pool as the filter condition.

Relationship between storage tiers and volumes

A volume that already belongs to a storage tier can be included in another storage tier. Figure 2-1:Relationship between storage tiers and volumes shows the relationship between storage tiers and volumes.

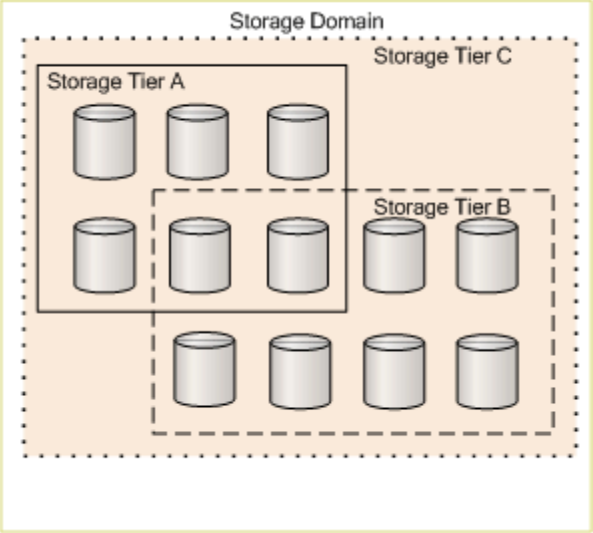

Figure 2-1: Relationship between storage tiers and volumes

You can specify the following filter conditions when creating a storage tier. For details about the values that can be specified for filter conditions, see the GUI Help.

Volume filter conditions:

- Device Number
- I/O Consumer
- Volume Status
- Subsystem
- Subsystem vendor
- Subsystem display model
- Subsystem serial number

- Ctrl. array group
- Array group
- Array Group Busy Rate
- Array Group Max Busy Rate
- RAID level
- Disk type
- Capacity
- Volume Lock Status
- Emulation type
- SLPR
- CLPR
- SYSPLEXID/DEVN
- VOLSER
- Logical group
- Port/HSD
- Disk RPM
- Disk capacity
- P-VOL’s Migration Group
- P-VOL’s MU Number
- ShadowImage
- TrueCopy Synchronous
- TrueCopy Asynchronous
- Universal Replicator
- Copy-on-Write Snapshot
- CVS
- Dynamic Provisioning
- Pool ID
- Encryption

Pool filter conditions:

- Free space
- Pool ID
- Over Provisioning Percent
- Over Provisioning Warning
- Over Provisioning Limit

Using the free space in a storage subsystem to create volumes

To enable effective use of a storage subsystem, you can search for specified sizes of contiguous free space and create volumes from the available capacity reported in the result. The required volume capacity can be specified in TB, GB, MB, or KB. *Contiguous free space* refers to areas of a volume physically next to each other that make up one, very large block of free space.

For Universal Storage Platform V/VM, when multiple areas of free space are found, all the areas are shown. For Universal Storage Platform V/VM or Hitachi USP storage subsystems, if there are multiple regions of free space, the capacity of each region will be displayed. For other storage subsystems, only the largest region of free space in the array group is displayed. For example, adjacent free space regions of 20 GB, 30 GB, and 60 GB are displayed as shown below.

- For Universal Storage Platform V/VM or Hitachi USP storage subsystems: 20, 30, 60
- For all other storage subsystems: 60

The number of volumes that can be created depends on the storage subsystem. When specifying free space in Universal Storage Platform V/VM, you can create a maximum of 2,048 volumes in one operation. When specifying free space in any other storage subsystem, you can create only one volume per operation. When specifying free space in a Hitachi storage subsystem, you can specify the emulation type from a drop-down menu.

Volumes can only be created in a storage subsystem that is registered in Device Manager.

To create volumes, create and execute a volume creation task. For details, see Volume creation tasks.

Operations to register storage subsystems with Device Manager, create volumes, and refresh storage subsystem information can be performed from Tiered Storage Manager.

If you are using the Hitachi AMS series, Hitachi WMS series, Thunder 9500V, or Thunder 9200, use port controller #0 to create volumes. If you want to use a port controller other than #0, change the settings for the `server.volumeCreation.param.defaultPortController` property in the `server.properties` file.

When creating a volume in an EVA SMI-S Enabled subsystem, the RAID level that is displayed in the free space list for the selected free space will be used for the volume. If necessary, use management software to assign a different RAID level to a volume.

Adding storage subsystems

Before a storage subsystem can be used with Tiered Storage Manager, it must be registered in Device Manager. You can register a storage subsystem in Device Manager by connecting to the Device Manager Web Client from the Tiered Storage Manager Web client.

For details about the settings necessary to add a storage subsystem in an IPv6 environment, see the Hitachi Device Manager Server Configuration and Operation Guide.

Note: Do not use an externally connected storage subsystem as a domain controller. If you migrate or shred the volumes in the external storage subsystem by mistake, the data might be lost. Volumes that have been locked by mistake are no longer accessible.

Refreshing storage subsystem information

After adding volumes or connecting external storage to a storage subsystem, refresh the information in the storage subsystem.

To refresh the storage subsystem information, in the Storage Domains subwindow, click the **Refresh Subsystem** button. The Refresh Device Manager Subsystem - Confirmation dialog box of Device Manager opens. From this dialog box, refresh the storage subsystem in Device Manager. If no storage subsystems have been registered in Device Manager, an error occurs.

Connecting external storage subsystems

When managing storage subsystems in tiers, you can map the volumes in external storage subsystems as Universal Storage Platform V/VM virtual volumes.

An external storage subsystem can be connected only if it meets all of the following conditions:

- Universal Storage Platform V/VM is registered as the Tiered Storage Manager storage domain.
- Universal Volume Manager is installed in the Universal Storage Platform V/VM subsystem.
- An external port is set up in the Universal Storage Platform V/VM subsystem and is connected to the target port of at least one external storage subsystem.
- An array group has been created in the external storage subsystem.

Note: If you want to establish an external connection to volumes in an SMI-S Enabled subsystem, be sure to use the same port number for the target port of the SMI-S Enabled subsystem as is being used by the current owner of the volumes in the SMI-S Enabled subsystem. If you specify different port numbers for these two items, you will be able to successfully set up an external connection, but you might not be able to access the external volumes. If the external volumes cannot be accessed, use management software to eliminate the problem between the internal volumes and the external SMI-S Enabled subsystem volumes.

Grouping target volumes

When you use Tiered Storage Manager, volumes that are to be migrated, locked, or unlocked, or volumes whose data is to be erased at the same time, must be grouped together. This group is called a *migration group*. Although multiple migration groups can be created in a single storage domain, each volume can belong to only one migration group. To prevent operating errors, you can use the characteristics of migration groups for the following operations:

- Grouping volumes you do not want to be selectable as migration sources

  You can use the group attributes to specify whether a group can be migrated. By adding the volumes you do not want to be selectable as migration sources to a group that cannot be migrated, you can prevent those volumes from being migrated by mistake.

- Grouping volumes you do not want to be selectable as migration targets

  Volumes that already belong to a migration group cannot be selected as migration targets. Therefore, by adding volumes you do not want to be selectable as migration targets to a migration group, you can prevent data from being migrated to those volumes by mistake.

  Tiered Storage Manager selects volumes without a path as migration target candidates. Therefore, if volumes without a path contain data that you do not want to be overwritten, register these volumes in a migration group.

Setting rules for migration

For each migration group, set the *LDEV selection rules,* which define the array groups from which Tiered Storage Manager selects the migration target volumes. There are two types of LDEV selection rules: the *array group avoidance rule* and the *array group selection rule*. For example, if you do not want an array group containing LDEVs that belong to another migration group to be selectable as a target array group, specify the array group avoidance rule. To ensure that the migration target is selected based on how array groups are used, specify the array group selection rule. The following describes an example of creating a migration group.

Registering a ShadowImage P-VOL and S-VOL in different migration groups

You should place a ShadowImage P-VOL and S-VOL in different array groups. If the P-VOL and S-VOL are placed in the same array group and an error occurs in that array group, the data on both the P-VOL and the S-VOL is no longer accessible, in which case creating volume copies becomes meaningless. Registering the P-VOL and S-VOL in different migration groups and then setting the array group avoidance rule prevents the P-VOL and S-VOL from being placed in the same array group by mistake.

To register a P-VOL and S-VOL in different migration groups:

1. Register the P-VOL in a migration group.

   To search for the P-VOL, use ShadowImage as the filter condition, and specify **P-VOL** as the value for the volume search.

2. Search for the S-VOL corresponding to the P-VOL.

   To search for the S-VOL, specify the name of the migration group to which the corresponding P-VOL belongs. Specify **P-VOL's Migration Group** or **P-VOL's MU Number** as the filter condition to perform volume search.
3. Register the S-VOL returned by the search in another migration group.
4. Specify the migration group containing the P-VOL in the array group avoidance rule for the migration group containing the S-VOL.
5. Specify the migration group containing the S-VOL in the array group avoidance rule for the migration group containing the P-VOL.

Adding a set of volumes used by a host to a migration group

By grouping the volumes that contain data used by a specific host, you can migrate host data all at one time. You might want to change array groups from which the migration target volumes are selected for the following reasons:

- You want to distribute volumes to as many array groups as possible because your application requires high-response performance.
- You want to select the migration target volumes from array groups with low busy rates and exclude array groups that have particularly high busy rates.
- You want to set aside as many array groups as possible that can be used as migration target candidates after migration so that the array groups are equally used.

You can specify array group avoidance rules and the array group selection rule at the same time. Tiered Storage Manager excludes the array groups related to the migration groups specified by array group avoidance rules, and then selects target LDEV candidates from the selected array groups based on the array group selection rule.

Array group avoidance rules

Use these rules to create a migration task in which the array groups containing the migration target LDEVs are different from the array groups containing the migration source LDEVs. Specifying an array group avoidance rule for a migration group is optional. When a migration task is created for a migration group that an avoidance rule has been specified for, Tiered Storage Manager will also exclude volumes in the avoided array groups from the migration target candidates.

Array group selection rule

Use this rule to specify how the array groups from which Tiered Storage Manager selects volumes are chosen. The three selection methods described below are available. The default method is `BalanceCapacity`.

BalanceCapacity

Specify this method when the equal use of array groups has priority over the location of the data. Tiered Storage Manager selects LDEVs so that the remaining capacity of each array group that can be selected as a migration target tends toward equality.

For example, you can use this method in the following case:

- When the storage subsystem has limited remaining storage capacity

  You can reserve more groups for use as migration target candidates for other migration groups.

MaximumCoverage

Specify this method when you want to distribute data among as many array groups as possible. Tiered Storage Manager selects LDEVs from as many array groups as possible.

For example, you can use this method in the following cases:

- When using an application that requires fast responses

  If all of the data is placed in the same array group, the required performance might not be achieved due to I/O conflicts. Distributing the data over multiple array groups would improve the response time.
- When using an application in which errors can be isolated by using partial blockage

  Distributing the data among multiple array groups can help prevent any disruption to system operations. If an error occurs in any array group, the data in the other array groups can still be accessed.

MinimumCoverage

Specify this method when you want to consolidate as much data as possible in a specific array group. Tiered Storage Manager selects LDEVs from as few array groups as possible.

For example, you can use this method in the following cases:

- When using an application that does not support error isolation by using partial blockage

  Distributing the data among different array groups increases the risk that an error in any one array group could halt system operation. Placing all the data in a highly reliable array group reduces the risk of hardware errors affecting the applications.
- When P-VOLs are distributed over a large number of array groups

  Distributing P-VOLs to many array groups reduces the number of array groups available for S-VOLs. Placing the P-VOLs and S-VOLs in a limited number of array groups increases the number of array groups that can be used, as well as simplifies the creation of copy pairs.

Figure 2-2:Differences in migration target volumes for each array group selection rule illustrates how Tiered Storage Manager selects LDEVs for each distribution method when six LDEVs are migrated. This example assumes that no array group avoidance rule is specified.

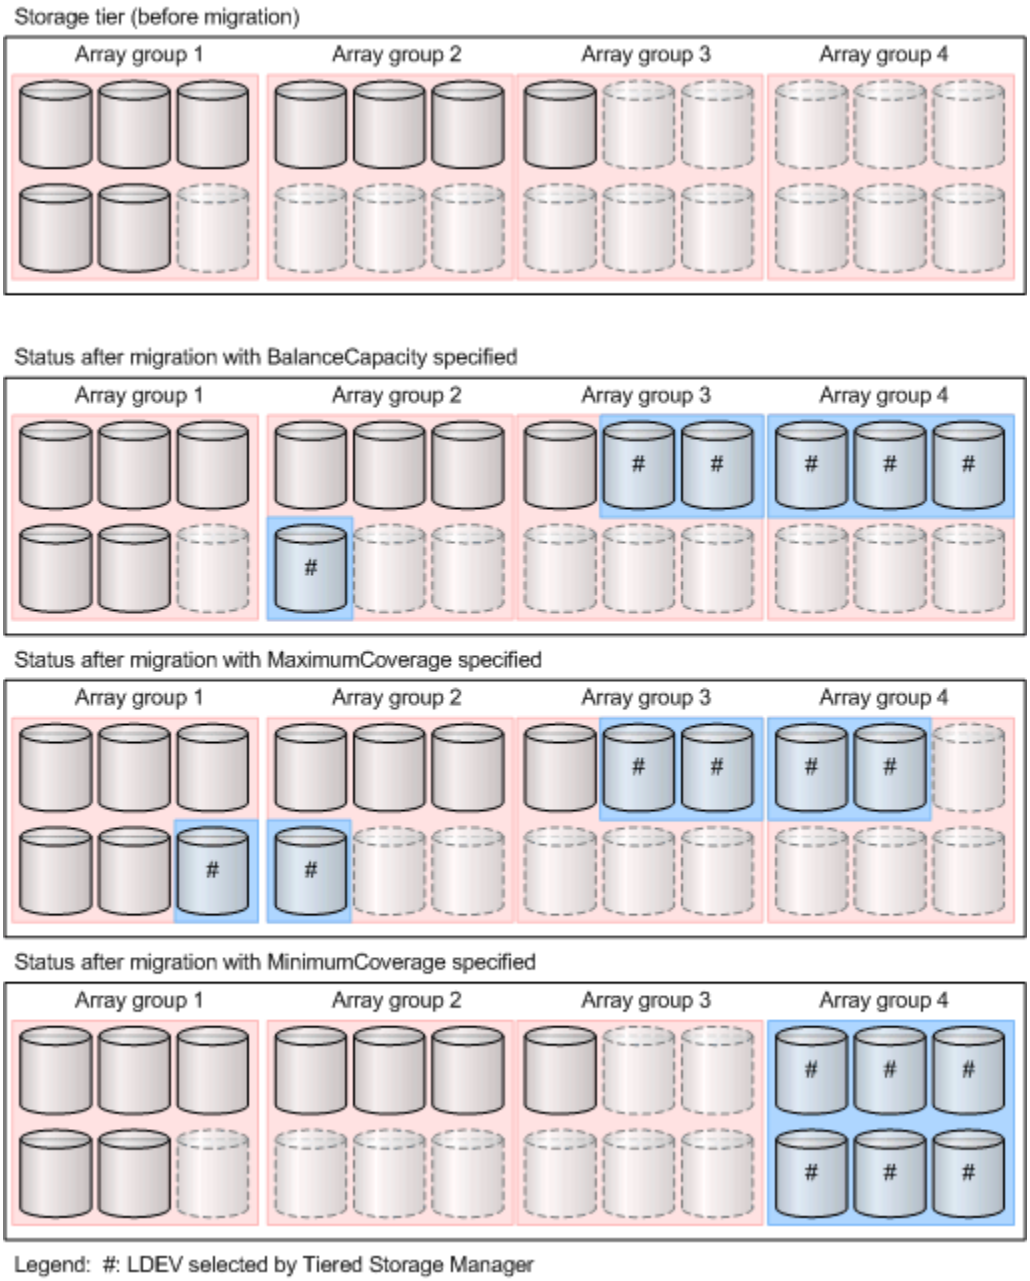

Figure 2-2: Differences in migration target volumes for each array group selection rule

For details about how LDEVs are selected based on the array group selection rules, see Details of the array group selection rules.

Preparing to migrate volumes

After you have defined migration groups and storage tiers, you can create *migration tasks*. Tiered Storage Manager displays all of the free volumes suitable as migration targets as migration target volume candidates. Select migration target volumes from the displayed target volume candidates. In addition to the migration target volumes you select, Tiered Storage

Manager displays volumes that can be used as migration target candidates. To perform migration, you can use the migration target volumes selected by Tiered Storage Manager or you can re-select migration target volumes from the volumes displayed as migration target candidates.

If the migration source volume belongs to a logical partition created by the SLPR function, Tiered Storage Manager chooses the migration target volume from the same logical partition. However, you can specify a migration target volume that belongs to an SLPR other than the migration source's SLPR. For details about SLPR, see the *Virtual Partition Manager User's Guide*. To prevent access from processes other than Tiered Storage Manager, the volumes selected as migration targets are reserved until migration is complete.

About Out of Execution volumes

If a task cannot be performed for a volume in a migration group, **Out Of Execution** is displayed as the status of that volume.

About intermediate volumes

Any intermediate volumes in a storage subsystem are displayed as open systems volumes in Device Manager. However, no intermediate volumes are displayed in Tiered Storage Manager.

If a storage tier has no free volumes with the same capacity as the migration source volume, normal volumes whose emulation type is OPEN-V that have a larger capacity than the migration source volume are also displayed as migration target candidates. When a migration target volume that has a larger capacity than the migration source volume is selected, the migration target volume with the larger capacity is deleted before migration and a new volume that has the same capacity as the migration source volume is created. When existing volumes must be deleted and new volumes created, the creation and deletion occurs before task execution. Therefore, the task execution takes more time than when no volumes must be created.

If there are multiple migration-target candidate volumes that have a larger capacity than the migration source volume, and all of these volumes are displayed, you must select the migration target volumes from a large list of volumes. By default, up to four migration target volumes are displayed. You can set how many migration target volumes that have a larger capacity are displayed by specifying the `server.migrationPlan.candidateCapacityGroupDisplayMaxCount` property in the `server.properties` file. For details about the `server.properties` file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

You cannot perform any operations that update the configuration of a storage subsystem while a volume is being created. Do not set a label or resource group for the volume until the migration task has finished.

Volume availability for migration

To use Tiered Storage Manager to migrate volume data, the migration source volumes and migration target volumes must satisfy the following conditions:

- The emulation types of the migration source and target volumes must be the same.

  Note, however, that if the microcode version of Universal Storage Platform V is 60-03-20-00/00-00 or later and both the migration source volume and migration target volume are DP volumes, or the migration source volume is a DP volume and the migration target volume is a DP pool, migration is possible between the emulation types OPEN-V and OPEN-0V.
- The number of blocks in the migration source and target volumes must be the same.
- The sizes of the migration source and target volumes must be the same.

  If the volumes are normal volumes whose emulation type is OPEN-V, migration is possible even if the migration target volume is larger than the migration source volume. In this case, the migration target volume is deleted, and then a volume that has the same capacity as the migration source volume is created.
- If the emulation type is not OPEN-V, both the migration source volume and migration target volume must be either CVS volumes or normal volumes.
- The volumes must have `LDKC:CU:LDEV` numbers in the range from `00:00:00` to `00:FF:FF`.
- If the volumes are configured as a TrueCopy Asynchronous, Hitachi TrueCopy Asynchronous for Mainframe, Universal Replicator, or Hitachi Universal Replicator for Mainframe® pair, the CLPR of the migration source and target volumes must be the same.

Volumes specifiable as migration targets when the migration source volume is a LUSE

When the migration source is a LUSE volume, Tiered Storage Manager uses the volumes selected as the migration target to automatically create a LUSE volume. Therefore, if the migration source includes a LUSE volume, you must select migration target volumes that satisfy all of the following conditions:

- The volumes must be usable as migration targets.

  To find out whether a volume can be used as a migration target, check the volume details. A volume can be a migration target if nothing appears in the **Migration Restrictions** field in the window.
- The migration target volumes must not satisfy any of the conditions listed in Table A-4:Volumes that can be used as migration sources but not as migration targets.
- Migration target volumes must be either all internal volumes or all external volumes.

- If the migration target volumes are external volumes, all the external volumes must have the same attributes (emulation type, I/O suppression mode, cache mode, and CLPR).
- The RAID level of all migration target volumes must be the same.
- The drive type of either all or none of the migration target volumes can be SATA.
- If you are migrating data to external volumes, all volumes should reside within the same external storage subsystem.

When the migration source volume is a LUSE volume in a Universal Storage Platform V/VM storage subsystem (microcode version: earlier than 60-05-12-00/00), do not issue any I/O activity to the volume while migration is being performed. If you do, data for the migration target volume might not be migrated correctly due to a defect in the microcode of the storage subsystem.

When the migration source volume is an external LUSE volume connected to a domain controller, the migration target volume must have a greater capacity than the migration source volume.

Filtering the number of candidate volumes

When creating a migration task, Tiered Storage Manager filters the volumes that can be used as migration target candidates. Filtering the number of candidate volumes prevents unnecessary volumes from being displayed and thus reduces memory usage. However, if the range of available LDEVs is predetermined, all the possible LDEVs might not be displayed as migration target candidates. If you want to display all available candidate volumes, use the `server.migrationPlan.candidateVolumeCountLimit` property of the `server.properties` file to disable filtering. For details about the `server.properties` file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

Example of filtering the number of candidate volumes

Table 2-1:Difference in the number of candidate volumes before and after filtering shows how Tiered Storage Manager filters the candidate volumes when the volume configuration in a migration group and storage tier is that shown in Figure 2-3:Configuration of the migration group and storage tier. This example assumes that the volumes are all OPEN-V and CVS.

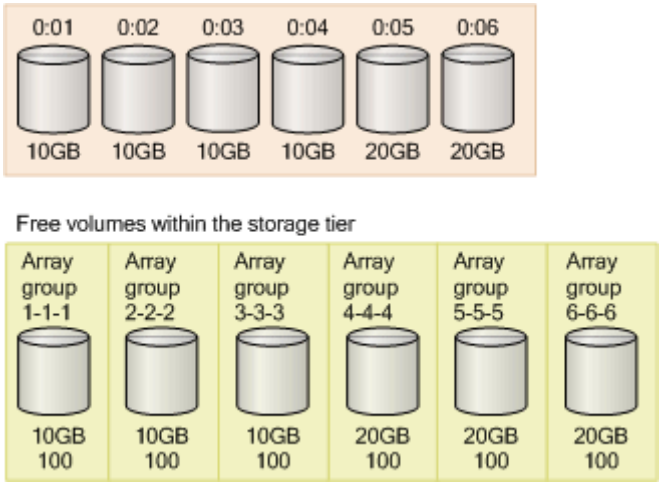

Figure 2-3: Configuration of the migration group and storage tier

The candidate volumes are filtered as shown in Table 2-1:Difference in the number of candidate volumes before and after filtering. For one migration source volume, one candidate volume is displayed in each array group. The filtering results in a reduction of candidate volumes from the initial 600 candidate volumes before filtering to 18 candidate volumes.

Table 2-1: Difference in the number of candidate volumes before and after filtering

Migration source volume	Classification of the candidate volumes	Number of candidate volumes before filtering	Number of candidate volumes after filtering
LDEV 0:01 LDEV 0:02 LDEV 0:03 LDEV 0:04	10 GB, array group 1-1-1	100	4
	10 GB, array group 2-2-2	100	4
	10 GB, array group 3-3-3	100	4
LDEV 0:05 LDEV 0:06	20 GB, array group 4-4-4	100	2
	20 GB, array group 5-5-5	100	2
	20 GB, array group 6-6-6	100	2

Migrating data to DP pools

Tiered Storage Manager allows DP pools to be specified as migration targets. You must create a storage tier by specifying a DP pool in advance. If a DP pool is specified as the migration target, Tiered Storage Manager dynamically creates DP pool volumes in the DP pool and migrates data to the volumes. You can specify a DP pool as the migration target if the following conditions are satisfied:

- The amount of free space in the migration target pool must be equal to or larger than the total size of the migration source volumes.
- The pool usage rate when DP pool volumes are created in the migration target pool must not exceed any of the threshold values specified for each DP pool.
- If a pool filter is specified for the target tier, select a pool that conforms to the pool selection rules.

If the migration source includes DP volumes, the following condition must also be satisfied:

- The same DP pool cannot contain both migration source and target volumes.

In the Web client, when a storage tier based on a DP pool is specified when creating a migration task, the initial status of the migration pair is `Do Not Migrate`. When the migration target DP pool is specified, the migration pair status changes to `Move`.

Using pool search conditions to create storage tiers

When using pool search conditions to create storage tiers for migration, Tiered Storage Manager version 6.3 and later supports a function where the target pool number is automatically set when the migration plan is created. For Tiered Storage Manager version 6.2 or earlier, you must manually set the target pool number when performing a migration.

Rules for excluding pools

If any pool in the storage tier meets any of the following conditions for excluding pools, that pool is excluded from the selection options in the migration plan.

Table 2-2: Conditions for excluding pools

#	Condition
1	The status of the pool is “Blocked”.
2	The usage rate of the pool exceeds either threshold 1 or threshold 2.
3	The usage rate of the pool will exceed either threshold 1 or threshold 2 if a migration is performed.
4	The maximum number of virtual VOLs (8,192) has already been created in the pool.
5	The Over Provisioning Percent of the pool exceeds the Over Provisioning Warning Threshold.
6	The Over Provisioning Percent of the pool will exceed the Over Provisioning Warning Threshold if a migration is performed.
7	The Over Provisioning Percent of the pool exceeds the Over Provisioning Limit Threshold.
8	The Over Provisioning Percent of the pool will exceed the Over Provisioning Limit Threshold if a migration is performed.
9	A migration source volume exists.

The target pool number is not returned if any of the conditions described in Table 2-3:Conditions where the target pool number is not returned are met.

Table 2-3: Conditions where the target pool number is not returned

#	Condition
1	The source volume capacity exceeds the upper limit for virtual VOLs (4 TB).
2	The source volume capacity is less than the lower limit for virtual VOLs (46 MB).
3	The source volume emulation type is not OPEN-V or OPEN-0V.
4	The source volume is a LUSE volume.
5	The source volume is in a pool within the storage tier.
6	The source volume is unavailable for migration.
7	The maximum number of volume groups (63,232) has already been created in the system.

Discarding zero-data after migrating to DP pools

While creating a migration task for a DP volume or a DP pool, you can select whether to release unused space on the target volumes after migration is performed to prevent unnecessary consumption of DP pools. You can check the status of the zero-data discard in the Tasks subwindow.

In the Web client, the check box **Discard zero data after migration** is selected by default. If the target volume is not a DP volume or DP pool, the zero-data discard option is disabled in the Web client.

If all of the following conditions are satisfied, the zero-data discard request is performed for the target volume after completion of the migration:

- The source volume is not a pair volume of ShadowImage or Copy-On-Write-Snapshot.
- The microcode version of Universal Storage Platform V/VM is 60-07-00-xx/xx or later.

Note: If the zero-data discard option is set to **Yes** and there are no volumes for which zero-data can be discarded, the request will be ignored. No error will be returned.

Erasing data after migration

While creating a migration task, you can select whether to erase the data on the migration source volumes after migration. To prevent unwanted disclosure of information, we recommend that you erase the data on the migration source volumes.

In the Web client, the check box **Erase remaining data on source volumes** is not selected by default. You can change the default to select this check box by modifying the `server.migration.dataErase.defaultValue` property in the

`server.properties` file. For details about the `server.properties` file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

When the migration source volume consists of DP volumes and Erase remaining data on source volumes is specified, the source volumes will be deleted after data erasure. In this case the targets for deletion depend on whether Open Volume Management is installed.

Creating and executing tasks to check the execution status

Storage subsystem operations in Tiered Storage Manager are created and executed as tasks. Tiered Storage Manager uses the following types of tasks:

- Migration tasks (can be created using the Web client or CLI client)
- Locking tasks (can be created using the CLI client)
- Unlocking tasks (can be created using the CLI client)
- Shredding tasks (can be created using the CLI client)
- Volume creation tasks (can be created using the Web client)
- External connection setup tasks (can be created using the Web client)

Tasks can be executed using the Web client or CLI client, and can be executed either immediately after they are created or at some later time. You cannot create an additional task for a migration group while a task for that migration group remains uncompleted.

The time required from task execution to completion varies depending on the volume. The estimated time from task execution to completion is displayed as a general guide for determining task execution order and timing, in the following cases:

- For migration tasks

  Estimated times are displayed for the migration time from task execution to completion, and immediately after migration until source volume data is erased and the optional zero-data discard task is completed.

- For shredding tasks

  The estimated time required from task execution to completion is displayed.

Estimated task times are calculated based on the previous migration and shredding times, and other data, in consideration of both the migration source and target volumes. When a task is executed for the first time, the estimated times can only be used as rough guides. As more tasks are performed, the estimated times become more accurate to the actual required times.

Note that the value used for calculating estimated task times was based on migrating a volume of 170 GB. When migrating a volume of less than 170 GB, the time actually required to complete the task might be significantly different than the estimated time.

You can check the status and progress of a task, including the status of tasks for individual volumes, in the Tasks subwindow or Task ID dialog box. Table 2-4:Task status describes the task status icons and their meaning.

Table 2-4: Task status

Task status icon	Description
	The task is pending execution. To execute a task that has this icon, select it, and then click the **Execute Tasks** button.
	The task is pending execution or is being executed. This icon is displayed if creation, deletion, or re-creation of a volume in a migration task is pending or is currently being performed. Wait until processing is completed.
	The requested processing has completed.
	*Stop* or *immediate stop* processing is being performed for the task.
	The task has been canceled or execution failed.

To check how much time has elapsed since a task has started, check **Time Elapsed**.

Event notification

Event notification is a useful function for managing tasks. This function allows you to specify that Tiered Storage Manager send an email when an event such as the completion of a task or the expiration of a specified period occurs. For example, if there is a volume you want to protect as read-only for three years starting one year after it was created, create a migration group `abcd` that contains the volume. Next, set a **Reminder Date** for the `abcd` migration group to one year later and specify settings so that an email will be sent indicating that the `abcd` migration group must be locked as read-only for three years. When this reminder email is delivered to you one year from now, you will remember to lock the volume.

Migration to another migration group

Because the volumes for which a task has finished are managed as a batch, volumes in a migration group can be migrated to another group after a locking task, unlocking task, or shredding task has been executed. A migration group for which a task has been created cannot be specified as a target migration group for a different migration task. Conversely, a task cannot be created for a migration group that has already been specified as a migration target.

Note: Errors, such as a failure to send email, cannot be verified from the Web client or CLI client. To ensure that no errors have occurred, regularly check the Tiered Storage Manager message log.

The following sections describe the various tasks that can be performed and the volumes for which tasks can be created. Notes on creating and executing tasks are also provided.

Migration tasks

When you execute a migration task, Tiered Storage Manager requests the storage subsystem to perform the migration. If a migration group contains volumes that cannot be migrated, only the volumes that can be migrated are migrated. In Tiered Storage Manager, the storage subsystem does not actually migrate all the volumes in a migration group immediately after the user executes a migration task. After the data on one migration source volume has been copied to a migration target volume, the storage subsystem exchanges the LDEV numbers of the migration source and target volumes. If multiple volumes are migrated at one time, the storage subsystem replaces the LDEV numbers in sequence after all of the data has been copied. Because the storage subsystem does not reassign any LDEV numbers until all data has been copied, no data is lost if an error occurs during copying.

When a request to stop a task is issued while the storage subsystem is performing migration, the LDEV numbers of some volumes are replaced before the task ends, but others are not. If zero-data discard is specified, you cannot stop a migration task while unused space is being released. You can specify either of the following modes to stop a task:

- Stop

  If the storage subsystem accepts a stop request from Tiered Storage Manager while data is being copied to a migration target volume, processing continues until the LDEV numbers are replaced.

- Immediate stop

  Processing is forcibly stopped even though the storage subsystem is copying data to the migration target volume.

After migration finishes, Tiered Storage Manager releases the volumes that were reserved while the migration was being performed. When the reservation status is released, Tiered Storage Manager changes the status of the migration task to `Success`. If an attempt to release the reservation status fails, the status of the migration task changes to `Failure` even though the migration was completed.

Locking and unlocking tasks

You can lock volumes in a migration group for a specified period of time. You can also specify write-protected or read/write-protected mode when locking volumes. The contents of read/write-protected volumes cannot be viewed. When the specified period expires, Tiered Storage Manager outputs an event notification, but it does not automatically unlock the volumes. After the specified period expires, use the CLI client to either extend the lock period or unlock the volumes.

To unlock volumes, you must create and then execute an unlocking task. A volume can only be unlocked after the specified lock period has expired. Note that this means that a lock can never be removed if you specified an

unlimited lock period when you created the locking task. In addition, when you execute a locking task for a migration group, the lock period will be applied to any previously locked volumes, as follows:

- If the lock period specified in the locking task is longer than the currently applied lock, the lock period of the previously locked volume will be extended.
- If the lock period specified in the locking task is shorter than the currently applied lock, the lock period of the previously locked volume will not be changed.

Tiered Storage Manager can lock volumes that satisfy all of the conditions below. If a volume does not satisfy all of the conditions, an error message is displayed when you execute the locking task.

- The volume is an open systems volume.
- The volume's emulation type is OPEN-3, OPEN-8, OPEN-9, OPEN-E, OPEN-K, OPEN-L, OPEN-V, or OPEN-0V.
- The volume is not an iSCSI volume.
- The volume is not a system disk volume.
- The volume is not a NAS user volume or a system volume.
- The volume is not a command device.
- The volume is not an on-demand volume.
- The volume is not a DP pool volume.
- The volume is not reserved by another program.
- If the I/O suppression mode is set, the volume is not mapped to an external LU.
- If VMA is set, the volume is not blocked.

If you are using a copy pair, all of the following conditions must be satisfied in addition to those above:

- The volume is not a Copy-On-Write Snapshot POOL or V-VOL.
- The volume is not a ShadowImage, TrueCopy Synchronous, or TrueCopy Asynchronous S-VOL.
- The volume is not a Universal Replicator S-VOL or JNL-VOL.

WARNING: You must also ensure the physical security of locked volumes. In particular, care is required when the locked volumes are located in an external storage connected to a Universal Storage Platform V/VM or the Hitachi USP. If the external storage subsystem is disconnected from the Universal Storage Platform V/VM or Hitachi USP, the volumes in the external storage subsystem will no longer be locked.

Locking volumes managed by Data Retention Utility

For volumes for which the Data Retention Utility attribute is not `Read/Write`, the status is displayed in Tiered Storage Manager as `Locked`. For volumes for which the Data Retention Utility attribute is `Read Only` or `Protect`, Tiered Storage Manager can either unlock the volumes or extend the lock period.

The Data Retention Utility has functionality called *expired lock*, which prevents unlocking even after the locking period has expired. If a migration group includes a volume that has an expired lock, an error occurs when you execute an unlocking task. If you want to unlock volumes whose lock has expired, disable the expired lock setting in Storage Navigator. For details about expired locks, see the *Data Retention Utility User's Guide*.

Shredding tasks

You can erase the data on all volumes in a migration group. If the migration group contains volumes whose data cannot be erased, such as locked volumes, only the data that can be erased is erased. You can specify either of the following shredding methods:

- ZERO-ONCE: Data is erased by overwriting the volume with 0s.
- DoD: Data is erased by overwriting the data multiple times by using a method that complies with the U.S. Department of Defense guidelines.

When data is erased after a migration initiated by a migration task, the ZERO-ONCE method is used. If you want to use the DoD method to erase data, create a shredding task.

You can create a shredding task if a volume in the migration group satisfies all the following conditions:

- A path for the volume has not been set.
- The Data Retention Utility attribute of the volume is `Read/Write`.
- The volume is not reserved by Volume Migration, TrueCopy Synchronous, TrueCopy Asynchronous, ShadowImage, Universal Replicator, or Copy-On-Write Snapshot.
- The volume is not a Volume Migration, TrueCopy Synchronous, TrueCopy Asynchronous, ShadowImage, Universal Replicator, or Copy-On-Write Snapshot paired volume.
- The volume is not a DP pool volume.
- The volume is not a system disk volume.

In addition, if the volume is a DP volume and a target migration group is not specified, the volume will be deleted after data is erased. Volume erasure operations vary depending on whether Open Volume Management is installed. For details, see Erasing data after migration.

Volume creation tasks

You can search for free capacity in a storage subsystem registered in Device Manager, and then create volumes in a specified array group. To do so, specify **Free Space(s)** as the search filter condition. Next, specify the number of volumes, volume size, and emulation type, and then create the volumes. This enables effective use of the free capacity available in storage subsystems.

Tiered Storage Manager automatically allocates LDEV numbers to the created volumes. You can create a Volume Creation Task when both of the following conditions are satisfied:

- The volume is an open systems volume.
- The storage subsystem is registered in Device Manager.

You can create only one volume creation task per array group. If there is already a volume creation task for the array group selected from the free space, an error will occur when the creation task is executed.

External mapping tasks

The volumes in an external storage subsystem registered in Device Manager can be connected to a storage domain. You can create a virtual volume by mapping the external storage subsystem as a single virtual array group. You cannot create either only a virtual array group or more than one virtual volume.

Tiered Storage Manager automatically allocates a name to the created virtual array group and an LDEV number to the virtual volume.

You can create an External Mapping Task if all the following conditions are satisfied:

- The storage domain is Universal Storage Platform V/VM.
- The external storage subsystem to be connected is registered in Device Manager.
- The emulation type of the virtual volumes is OPEN-V.
- A storage path has been established between the external storage subsystem to be connected and the Universal Storage Platform V/VM.

Do not access volumes in the external storage subsystem from a host connected to that subsystem. Also, do not use the copy function of the external storage subsystem to access the volumes. If volumes have been connected as storage domain volumes, you must access the volumes from the storage domain.

Notes on tasks

- From the time you create a task until the task ends, do not change the configuration of the volumes included in the task. If you change the volume configuration, the task might fail or produce unexpected results. For example, if you use Storage Navigator or Device Manager to change a migration source volume to a LUSE after you have created a migration task, it is possible that data on only some LDEVs on the LUSE volume will be migrated. This might result in the LUSE volume being incomplete.
- If you perform any of the following operations for the migration source volume, migration will be interrupted:
  - A compatible XRC operation
  - A concurrent Copy (CC) operation
  - A Hitachi TrueCopy for Mainframe or TrueCopy operation that changes the volume status to a status other than Suspended

    - A Hitachi Universal Replicator for Mainframe or Universal Replicator operation that changes the volume status to COPY
    - A Hitachi ShadowImage for Mainframe or ShadowImage operation that results in the following volume status:

      Mainframe volume pair: SP-Pend or V-Split

      Open systems volume pair: COPY(SP) or PSUS(SP)
    - Universal Replicator and Hitachi Universal Replicator for Mainframe operations
- Do not set Cache Residency on the volume being migrated.
- If an error occurs in a storage subsystem during task execution, Tiered Storage Manager might automatically retry the task until the cause of the error is removed. This is not a serious problem if the cause of the error is removed within a few minutes. However, if the error persists, the task does not terminate. If completion of a task is taking an excessive amount of time, use the Web client or CLI client to check the task details. If an error message has been output for the task, take action as indicated in the message.
- An error occurs if you attempt to create a new task in a storage domain in which a migration has failed. Create the task after refreshing the storage domain.
- If you perform migration while access from hosts is stressing the storage subsystem, the migration might fail. If migration fails, temporarily reduce the storage subsystem load and then retry migration.
- If you frequently perform update I/Os for volumes being migrated, copying for Volume Migration, ShadowImage, Hitachi ShadowImage for Mainframe, and Compatible FlashCopy will slow down, because the copying of differential data is being repeatedly executed. Table 2-5:Capacity of migrated volumes and estimated delay in copy time lists how much the copying slows down when update I/Os for a migration source volume occur 50 times per second. Although the copy speed of Volume Migration, ShadowImage, Hitachi ShadowImage for Mainframe, and Compatible FlashCopy depends on the number of pairs, at a maximum, copy time doubles.

Table 2-5: Capacity of migrated volumes and estimated delay in copy time

Capacity of migrated volume (MB)	Copying speed delay (minutes)
0 - 1,000	4
1,001 - 5,000	18
5,001 - 10,000	37
10,001 - 50,000	186
50,001 - 100,000	372
100,001 - 500,000	1,860
500,001 - 1,000,000	3,720
1,000,001 - 2,150,400	9,667

- If the message KATS50335-E appears during migration, a target migration LDEV might not have been formatted. Check the LDEVs. If there are any unformatted LDEVs, format them. After formatting, refresh both Device Manager and Tiered Storage Manager, and then perform the migration again. If there are no unformatted LDEVs, see the Hitachi Tiered Storage Manager Software Messages to take appropriate action.
- A migration task being executed during storage subsystem maintenance might fail. Do not execute tasks while a storage subsystem cache or drive is being added, replaced, or removed or while a microprogram is being replaced.
- If you update configuration information on an SMI-S Enabled subsystem, an error might occur when a volume creation task or external connection setup task is executed because it takes a long time to reflect the information to the SMI-S provider. If the message KATS62010-E is output, see the Hitachi Tiered Storage Manager Software Messages to take appropriate action. If you still cannot resolve the problem, perform the following procedure:

  1. Update the SMI-S provider information.
  2. Refresh Device Manager.
  3. Refresh Tiered Storage Manager.
  4. Execute the task.

  You must also perform the above procedure if the message outputs the following:

  ```
  An error was detected in Device Manager. code: "7208"
  contents: An attempt to create an external group has failed
  because the combination of the specified port "port-name"
  external storage subsystem volume (LUN = "0FFFFFFFF(-1)") is
  invalid.
  ```

Updating the Tiered Storage Manager database

Tiered Storage Manager, its prerequisite product Device Manager, and the related product Tuning Manager have separate databases that contain storage subsystem information. Updating the database of a Hitachi Storage Command Suite product is referred to as *refreshing*. Note that you cannot create or execute a task in Tiered Storage Manager while the Tiered Storage Manager database is being refreshed. While the database is being refreshed, you can only display a list of storage domains, view properties, and collect task information.

Refreshing the database affects the performance of Device Manager and Tuning Manager. Refresh the database only if a warning message appears in the Tiered Storage Manager window. Obtaining the information required for updating the database is, in itself, a time-consuming process.

Purpose of refreshing

There are two reasons for refreshing a database:

- To refresh the array group busy rate

  You can collect information about the array group busy rate from Tuning Manager.

- To apply a configuration change in a storage subsystem to the Tiered Storage Manager database

  When you use Device Manager to create an LDEV in a storage subsystem or set the path to a host for an LDEV, the configuration change in the storage subsystem is applied to the Device Manager database. However, you need to refresh the Tiered Storage Manager database in order to apply this change to this database.

  Tiered Storage Manager monitors the Device Manager database at user-specified intervals to check whether the latter has been updated. When it detects that the Device Manager database has been updated, Tiered Storage Manager displays a warning message. Make sure that you refresh the Tiered Storage Manager database when a warning message appears. Note, however, that if only the volume label information is updated in the Device Manager database, Tiered Storage Manager determines that no refreshing is required and does not display a warning message.

  To specify how often to detect updates of the Device Manager database, use `server.repository.dvmModifyCheck.pollingIntervalInMinute` in the `server.properties` file.

  If you use software other than Device Manager, such as Storage Navigator, to change the configuration of a storage subsystem, the Device Manager database matches the Tiered Storage Manager database, but neither matches the actual storage subsystem information. In such cases, refresh the Device Manager database, and then the Tiered Storage Manager database.

Automatic refreshing

You can specify automatic refreshing for Tiered Storage Manager at the following times:

- When the Tiered Storage Manager server is started
- Once each day at a specified time
- When a refresh of the Device Manager database is detected

Depending on the type of automatic refresh, there will be the following time lag until the Device Manager database and the Tiered Storage Manager database match:

- Until the next time the Tiered Storage Manager server is started
- Until the specified time on the next day (maximum of 24 hours)
- Until the next time the last update time of the Device Manager database is checked (30 minutes to 24 hours, depending on the specified interval)

To specify options for automatically refreshing the Tiered Storage Manager database, set the `server.repository.autoRefresh.pollingIntervalInMinute`, `server.repository.autoRefresh.serverStarted`, and

`server.repository.autoRefresh.recurringLocalTime` properties in the `server.properties` file. For details about these properties, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

When automatic refreshing is enabled in Replication Manager that is installed on the same server as Device Manager, Replication Manager periodically refreshes the Device Manager database. Tiered Storage Manager is refreshed whenever Device Manager is refreshed. Therefore, depending on whether automatic refreshing is enabled in Replication Manager, the information in Device Manager and Tiered Storage Manager will be constantly updated even if the configuration of Device Manager remains unchanged. To prevent the Tiered Storage Manager database from being updated numerous times every day, in the `server.properties` file, turn off automatic refreshing of the Device Manager database. If you do not want a warning to be displayed, also disable the detection of updates in the Device Manager database.

Refresh status indicators

Before using Tiered Storage Manager, check whether the storage domains have been refreshed. Use the Storage Domains subwindow or a *storage-domain-name* subwindow to check the refresh status. There are six refresh statuses:

- Normal (refresh succeeded)

  The information was refreshed successfully. The database and cache are correct.

- Unexecuted

  When a storage domain is created, information is automatically refreshed. However, there might be a time lag from immediately after a storage domain is created until the information is refreshed. This status is in effect during this period. If this automatic refresh fails, you should refresh the information manually.

- Currently refreshing

  Information is currently being refreshed.

- Cache failure (refresh failed)

  Refresh processing failed because of an error. The status of both the database and the cache is abnormal.

- Incomplete (collection of Tuning Manager performance information failed)

  Only the collection of performance-related information from Tuning Manager failed. In this case, you need to refresh the information again only if Tiered Storage Manager was unable to collect information from an array group containing a volume for which you wanted to collect performance information.

- Refresh required

The storage domain information is not up to date. Refresh the storage domain.

Note: If the message KATS51407-E is output, refresh the storage domain again. The refresh status may be shown as **Normal** even though the refresh operation was unsuccessful.

Searching volumes, pools, and free space

This section describes the filter conditions for volumes, pools, and free space when used in the following cases:

- Creating a storage tier in a storage domain
- Choosing volumes in a storage tier to be migration targets
- Obtaining a list of volumes in a storage domain or storage tier (only volume filter conditions can be used)
- Creating volumes in the free capacity of a storage subsystem registered in Device Manager.

For any one operation, you can specify only one type of volume filter condition: volumes, pools, or free space.

Condition elements of volume filter conditions

Volume filter conditions are based on the attributes of volumes in a storage subsystem. You can define from 1 to 30 conditions. The following filtering operations can be performed:

- Filtering based on a single condition
- Filtering that uses AND to combine multiple filter conditions
- Filtering that uses OR to combine multiple filter conditions

For details about volume filter conditions, see Creating a storage tier by specifying volume filter conditions or Creating a volume group by filtering volumes.

Searching for volumes in a storage domain

You can specify the following items as the condition elements in the filter conditions used for volume searches. For details about the values that can be specified for filter conditions, see the GUI Help.

- Device Number
- I/O Consumer
- Volume Status
- Subsystem
- Subsystem Vendor
- Subsystem Display Model
- Subsystem Serial Number
- Ctrl. Array Group
- Array Group

- Array Group Busy Rate
- Array Group Max Busy Rate
- RAID Level
- Disk Type
- Capacity
- Volume Lock Status
- Emulation Type
- SLPR
- CLPR
- SYSPLEXID/DEVN
- VOLSER
- Logical Group
- Port/HSD
- Disk RPM
- Disk Capacity
- P-VOL's Migration Group
- P-VOL's MU Number
- ShadowImage
- TrueCopy Synchronous
- TrueCopy Asynchronous
- Universal Replicator
- Copy-On-Write Snapshot
- CVS
- Dynamic Provisioning
- Consumed Capacity
- Consumed Capacity Percentage
- Pool ID
- Label
- Encryption

- Generally, the range of values displayed in a drop-down list when a filter condition is selected is limited to values in the storage domain. As exceptions to this general rule, the following values are also displayed in drop-down lists:

  - A value not in the storage domain and set as a filter condition by using the CLI

  - A previously set value that is no longer applicable because of a change in the storage configuration
- If a storage subsystem is not registered or cannot be registered with Device Manager, the subsystem name might be displayed more than once in the subsystem drop-down list. Choosing any of the identical subsystem names will result in the same search results.

Searching for volumes in a storage subsystem registered in Device Manager

Searches for volumes in a storage subsystem registered in Device Manager are possible only from the Web client. You can specify the following condition elements in the filter conditions:

- Device Number
- Label
- Capacity
- Path
- Array Group
- RAID Level
- Volume Status

Condition elements of pool filter conditions

You can specify pool filter conditions based on the attributes of the storage subsystem pools. You can only use a pool filter condition to specify pools that are managed in a storage tier. The following filtering operations can be performed:

- Filtering based on a single condition
- Filtering that uses AND to combine multiple filter conditions
- Filtering that uses OR to combine multiple filter conditions

You can specify the following items as the condition elements in the filter conditions for pool searches. For details about the values that can be specified for filter conditions, see the GUI Help.

- Free Capacity
- Pool ID
- Over Provisioning Percent
- Over Provisioning Warning
- Over Provisioning Limit

Condition elements of free space filter conditions

You can specify one or more filter conditions based on the attributes of the free space in a storage subsystem registered in Device Manager. The following filtering operations can be performed:

- Filtering based on a single condition
- Filtering that uses AND to combine multiple filter conditions
- Filtering that uses OR to combine multiple filter conditions

You can specify the following items in the filter conditions. For details about the values that can be specified for filter conditions, see the GUI Help.

- Array Group
- Capacity
- Total Free Capacity

- RAID Level
- Emulation Type
- Disk Type
- Disk RPM
- Disk Capacity
- Subsystem Vendor
- Subsystem

The condition elements that are displayed depend on the storage subsystem type.

Depending on the SMI-S provider, if a storage pool of an SMI-S Enabled subsystem does not contain any volume, a free space search might result in displaying no information in the search result. If this happens, use management software to create a volume. After creating a volume if no information is still displayed, update the information about the SMI-S provider, and then refresh Device Manager.

Notes on filter conditions

- If `server.checkOutVolumeRange=false` is specified in the `server.properties` file for Tiered Storage Manager, one-byte spaces, upper-case characters, and lower-case characters are used exactly as entered in the volume filter conditions. For details about the `server.properties` file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.
- If you want to specify `VOLSER` or `SYSPLEXID/DEVN` as a filter condition, you must install Mainframe Agent on the mainframe host and register it in Device Manager. For details about installation and registration, see the *Hitachi Storage Command Suite Server Installation Guide* and *Hitachi Device Manager Mainframe Agent User's Guide*.
- If you use Device Manager to add or refresh a storage subsystem, it might take a while for the results to be reflected in Tiered Storage Manager. If the expected search results are not displayed, even if you perform a search without specifying filter conditions, wait a while, and then try again.

Allocating host paths to volumes

You can allocate and release host paths for volumes. You can also allocate and release paths for the volumes in a migration group or storage tier, as well as for a maximum of 100 volumes at one time. If you attempt to simultaneously use volumes that already have a path allocated and those that do not, an error occurs.

A path can be allocated to a volume that satisfies all of the following conditions:

- The volume is an open systems volume.
- The volume is not specified as a system disk.
- The volume is not an On-Demand volume.

- The volume is not a DP pool volume.
- The volume is not a Copy-On-Write Snapshot pool volume.
- The volume is not a Universal Replicator journal volume.
- The volume is not a Hitachi Universal Replicator for Mainframe journal volume.
- The volume is not reserved as a migration target.
- The volume is not included in a standby or running task.

If an attempt to allocate or release a path fails, refresh the Device Manager and Tiered Storage Manager databases.

To allocate a path to a volume, you can create a LUSE volume consisting of selected volumes. When you use **Volume List** of the migration group to view the volumes in the LUSE volume, LDEVs other than the representative LDEVs will no longer be displayed.

Using performance information acquired from Tuning Manager

It is necessary to perform migration in order to improve performance of volumes with performance problems. Enabling connection to Tuning Manager allows you to check performance information.

Connecting to Tuning Manager

We recommend that you check volume performance information before performing migration in order to ensure efficient migration. You can connect to Tuning Manager from the Tiered Storage Manager Web client to display the Tuning Manager Historical Report dialog box. You can display the Historical Report dialog box for each volume, migration group, or storage tier.

You can connect to Tuning Manager if all of the following conditions are satisfied:

- You have View permission for Tuning Manager.
- Volumes have been allocated to the host (application).
- Tuning Manager server 6.2, or later, has been installed.
- A Tuning Manager agent has been installed on the storage subsystem that is being used to acquire performance information.

Acquiring array group performance information

You can prevent access from being concentrated on a specific array group by migrating data on the volumes in array groups with high usage rates to array groups with low usage rates. To do this, check the *average busy rate* (**Busy %**) and *maximum busy rate* (**Max Busy %**) of the array group, and then select the volume to be migrated. Displaying the array group busy rate in a Tiered Storage Manager window requires that you install Tuning Manager. Tuning Manager periodically collects performance information from the storage subsystems being monitored. For some storage

subsystems, however, the busy rate might not be displayed even if Tuning Manager has been installed. For details about the storage subsystems that Tuning Manager can monitor, see the *Hitachi Tuning Manager Hardware Reports Reference*.

Tiered Storage Manager can acquire the weekly or monthly busy rate from Tuning Manager when a storage domain is refreshed. Use the properties file of the Tiered Storage Manager server to specify whether to acquire the busy rate and whether to acquire the busy rate on a weekly basis or a monthly basis. For details about the properties file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

Using mainframe volumes

Mainframe volumes used in OS/390 or z/OS can be migrated and shredded by Tiered Storage Manager. If you attempt to use Tiered Storage Manager to manipulate a volume used in another mainframe OS, the mainframe system might no longer be able to access the volume.

You can perform the following operations for mainframe volumes:

- Volume management
  - Display volume information
  - Search volumes
  - Add volumes to a migration group
  - Delete volumes from a migration group
  - Create storage tiers
- Migration
- Shredding
- Event notification

Note: Tiered Storage Manager cannot check some of the conditions used to determine whether a mainframe volume can be migrated. Consequently, migrating a mainframe volume might cause the following problems:

- If the selected volume is used on a host on which Mainframe Agent is not installed, or if Mainframe Agent is installed but does not collect information from that volume, any data on the migration target volume will be completely erased.
- An error might occur during migration to or from a copy pair volume.

To prevent the above problems, make sure that Mainframe Agent collects information about all mainframe volumes.

To specify that information is collected about all mainframe volumes:

1. Install Mainframe Agent on all the hosts.
2. Use the Mainframe Agent initialization parameter `DEVN` to specify the range of device numbers for all volumes.

Note that the device numbers specified for collecting information are used by the mainframe, and are not the same as the device numbers used by Tiered Storage Manager.

3. Register volumes that cannot be used with Tiered Storage Manager in a dummy migration group.

   When registering this migration group, set **No** for **Can Migrate** in the migration group attributes.

To check the information about volume pairs, perform either of the following:

- Display the management window of each storage product using Storage Navigator, and check the status of each volume based on the LDEV number.
- Specify `DEVN` as a search key in Business Continuity Manager, and then obtain the pairing status of each volume.

For details about Mainframe Agent, see the *Hitachi Device Manager Mainframe Agent User's Guide*. For details about Storage Navigator, see the *Storage Navigator User's Guide*. For details about Business Continuity Manager, see the *Hitachi Business Continuity Manager User's Guide*.

3

Using the Web client

This chapter explains how to log in to the Tiered Storage Manager Web client, register users, and migrate volumes.

- About the Tiered Storage Manager Web client
- Login and logout
- Registering licenses
- Setting up user accounts
- Setting up storage domains
- Setting up storage tiers
- Creating volumes
- Adding and refreshing storage subsystems
- Connecting external storage subsystems to a storage domain
- Setting up migration groups
- Performing migration
- Working with tasks

About the Tiered Storage Manager Web client

This section explains the structure of the windows of the Tiered Storage Manager Web client and notes on operations.

Structure of the windows of the Tiered Storage Manager Web client

Figure 3-1:Structure of windows of the Tiered Storage Manager Web client shows the structure of the windows of the Tiered Storage Manager Web client.

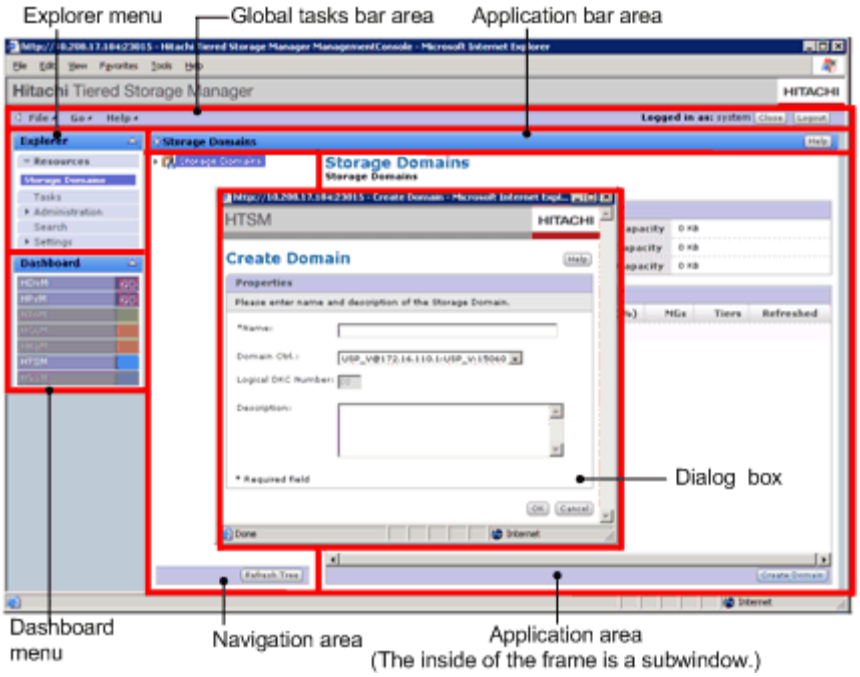

Figure 3-1: Structure of windows of the Tiered Storage Manager Web client

Global tasks bar area

Displays the function menus and action buttons that are used with Tiered Storage Manager, as well as information about the logged-in user.

Explorer menu

Displays the Tiered Storage Manager menu items. Choosing a menu item displays the corresponding information in the navigation area and the application area.

Dashboard menu

Displays a list of Hitachi Storage Command Suite products. An enabled product can be started by clicking its **Go** button.

Application bar area

Displays the menu item and action buttons corresponding to the item chosen from the **Explorer** menu. This area also contains an icon for hiding or displaying the navigation area.

Navigation area

Displays in tree format the objects that belong to the menu item chosen from the **Explorer** menu. When you expand the object tree and choose an object, information about the selected object appears in the application area.

Application area

Displays information about the object that was chosen from the **Explorer** menu and object tree.

Subwindow

A display in the application area is called the *object-name* subwindow.

Dialog box

A pop-up window that is displayed when a button or icon is clicked is called the *button-name-or-icon-name* dialog box.

Notes on GUI operations

Do not perform the following actions when using the Tiered Storage Manager GUI:

- Click browser buttons such as **Next** or **Back**
- Select an item from the right-click menu in a window
- Use shortcut keys, such as **Back Space** or **Tab**
- Select menu bar items in the browser window

These actions will not damage the database that is used internally by Tiered Storage Manager. If you do happen to perform one of the above actions, update the window by using the browser.

Notes on Internet Explorer 7.0/8.0

- When the tab function is used in Internet Explorer 7.0/8.0, resizing a dialog box might cause the contents to overlap.
- The menu and address bars might not be displayed or hidden correctly.

Notes on Windows Vista and Windows 7

- If 4-byte characters are used, operation is not guaranteed.

Login and logout

Tiered Storage Manager is started from its User Login dialog box.

You can also start Tiered Storage Manager from another Hitachi Storage Command Suite products's Web client (such as, Device Manager or Tuning Manager) by choosing the Tiered Storage Manager **GO** link in the **Dashboard** menu. This method does not require entry of a user ID or password.

Logging in

To start Tiered Storage Manager from its User Login dialog box:

1. Enter the login URL below in the Web browser's address bar.

   Be sure to bookmark this page.

   `http://{`*Tiered-Storage-Manager-server-name*`|`*IP-address*`}:`*port-number-of-HBase-Storage-Mgmt-Web-Service-in-Tiered-Storage-Manager-server*`/TieredStorageManager/`

   For IPv6 addresses, a connection cannot be established unless the IP address is enclosed in brackets (`[]`).

   Example of an IPv6 address:

   ```
   http://[1080:2C14:D30:BA04:275:806:270C:418A]:23015/
   TieredStorageManager/
   ```

   Only global addresses can be specified, not site-local addresses or link-local addresses.

   If using Internet Explorer 6.0 as your Web browser in an IPv6 environment, you cannot access the Tiered Storage Manager server by specifying an IPv6 address. Instead, specify the IPv6 address and host name of the Tiered Storage Manager server in the `hosts` file, and then access Tiered Storage Manager by using the host name.

   The User Login window appears.

2. Enter your user ID and password and then click the **Login** button.

   After installation, use the following user ID and password to log into Tiered Storage Manager for the first time:

   - User ID: `system`
   - Password: `manager`

   If you have enabled authentication using an external authentication server, use the password registered in that server.

   Tiered Storage Manager Web client's initial window appears. If you enter the wrong password repeatedly, you will be locked out. If this happens, have a user with User Management permission unlock the account. For details on how to unlock accounts, see Changing the lock status of a user.

Logging out

To log out, click the **Logout** button in the global task bar area. The Tiered Storage Manager Web client will log out the user and then terminate.

Once you are logged out, you can no longer use Hitachi Storage Command Suite products that were started from the terminated Web client's **Dashboard** menu. If you click the **Close** button instead of logging out, the browser closes, but you can continue to use other Hitachi Storage Command Suite products that were launched from the **Dashboard** menu.

Registering licenses

Licenses are necessary for using Tiered Storage Manager. If a license has not been registered for a monitored subsystem, a warning is displayed in the global tasks bar area. In this case, you need to register a license.

To register or update a Tiered Storage Manager license:

1. Display the License Information dialog box by using one of the following methods:
   - Click the **License** button in the Tiered Storage Manager login dialog box.
   - After logging in to Device Manager, in the global task bar area, select **Help**, and then **About**.
   - After logging in to Tiered Storage Manager, in the **Explorer** menu, select **Settings**, and then **License**. In the License Information dialog box, click the **Edit License** button.

   The currently registered license information is displayed.
2. Specify the license key or license file.

   To directly enter the license key, select the **Key** radio button, and then enter the license key.

   To use a license file, select the **File** radio button, click **Browse**, and then select the license key file.
3. Click the **Save** button.

Setting up user accounts

In Tiered Storage Manager you grant permissions to users individually so that the actions that users can perform are limited to match user goals and to prevent inappropriate user operations.

Tiered Storage Manager comes with the following users already set up:

- `System` account

  The built-in account for all Hitachi Storage Command Suite products. This account can use the functionality of all Hitachi Storage Command Suite products, including Tiered Storage Manager. The default password for the `System` account is `manager`.
- `HaUser` account

The default user account used by Device Manager agents.

User permissions

Tiered Storage Manager includes the following user permissions. Multiple permissions can be set for a single user.

- Admin permission

  Allows you to manage all storage domains and view all Tiered Storage Manager resources.

- Modify permission

  Allows you to use all functions except for those used to create and delete storage domains.

- Execute permission

  Allows you to execute tasks and view all Tiered Storage Manager resources. It is used to execute previously created tasks such as those for nightly batch processing.

- View permission

  Allows you to view all Tiered Storage Manager resources. The `View` permission is set automatically when you set the `Admin`, `Modify`, or `Execute` permission.

- User Management (Admin) permission

  Allows you to manage all users, regardless of any other permissions you have been granted. The Admin permission for User Management is hereafter called the User Management permission.

Operations that can be performed for storage domains, storage subsystems, migration groups, storage tiers, and tasks based on user permissions are listed in Table 3-1:Operations permitted for Tiered Storage Manager resources based on user permissions.

Table 3-1: Operations permitted for Tiered Storage Manager resources based on user permissions

Category	Operations	Admin	Modify	Execute	View
Storage Domain	Create or delete	Y	--	--	--
	Edit attribute information	Y	Y	--	--
	View	Y	Y	Y	Y
	Refresh	Y	Y	--	--

Category	Operations	Admin	Modify	Execute	View
Storage Tier	Create, delete, or edit attribute information	--	Y	--	--
	View	Y	Y	Y	Y
	Assign or release storage	--	Y#1	--	--
	Display volume and array group performance information	Y#2	Y#2	Y#2	Y#2
Storage subsystem	Search for free capacity	Y	Y	Y	Y
	Create or externally connect volumes	--	Y	--	--
	Register storage subsystem in Device Manager or perform a refresh	--	Y#1	--	--
Migration Group	Create, delete, or edit attribute information	--	Y	--	--
	Add or remove volumes	--	Y	--	--
	View	Y	Y	Y	Y
	Assign or release storage	--	Y#1	--	--
	Display volume and array group performance information	Y#2	Y#2	Y#2	Y#2
Task	Create, delete, change, cancel, or stop	--	Y	--	--
	Execute	--	Y	Y	--
	View	Y	Y	Y	Y
	Display the estimated time	--	Y	Y	--
	Display volume and array group performance information	Y#2	Y#2	Y#2	Y#2

Category	Operations	Admin	Modify	Execute	View
Legend: Y: Operation permitted --: Operation prohibited #1 You must also have the Modify permission for Device Manager to perform this operation. `All Resources` of Device Manager must be assigned. #2 The View permission for Tuning Manager is required.					

As an exception to the above table, users with either the Execute or View permission can create a migration plan by using the CLI.

User Management permission is required to perform the following operations:

- Adding and deleting users
- Setting user permissions
- Viewing and editing other users profiles
- Changing other users passwords
- Setting password conditions
- Setting up the automatic locking of user accounts
- Locking and unlocking user accounts
- Changing the authentication method
- Setting up warning banner messages

The User Management permission is not required to perform the following operations:

- Viewing and editing your own user profile
- Changing your own password

Adding a user

When logging in to Tiered Storage Manager for the first time after installation, log in using the `System` account, and then add users who have User Management permissions or Tiered Storage Manager Admin permissions as necessary.

When linking to an external authentication server is used, the default authentication method for added users is external authentication.

You do not need to add users who log in by linking to an external authorization server, because those user accounts are managed by authorization group.

To create a new user account:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**. In the navigation area, select **Users**.

   The Users subwindow appears. The **User List** appears in the Users subwindow.

2. Click the **Add User** button in the Users subwindow.

   The Add User dialog box is displayed.

3. Enter the information for the new user, and then click the **OK** button.

   For details about specifiable characters, see the online Help.

   The following characters have been added as usable characters in Hitachi Storage Command Suite products version 6.3:

   ! $ % & ' ( ) * = \ ^ |

   If you use a user ID that includes any of these characters, you cannot log in to Replication Manager version 6.2 or earlier if that version is installed on the same computer as Tiered Storage Manager.

   If you register the following users, specify a character string that includes a realm:

   - A user that is authenticated by another RADIUS server connected via a RADIUS server that is specified as the connection destination in the Tiered Storage Manager server properties file.
   - A user that is registered in a Kerberos server and belongs to a realm other than the realm specified as the default in the Tiered Storage Manager server properties file.

   A password is optional if linkage with an external authentication server is enabled in the Tiered Storage Manager server properties file.

   Passwords must satisfy any set password conditions, such as the required character combinations or minimum number of characters. For details about the password conditions, see Setting password conditions.

Setting permissions for a user

To set Tiered Storage Manager permissions for other users, log in as a user with User Management permissions.

You cannot set permissions for users who log in by using an external authorization server because their accounts have not been registered in the Hitachi Storage Command Suite products. For details about how to set up permissions for users who use an external authorization server, see Enabling external authentication using an external authorization server.

Keep in mind the following conditions while setting permissions:

- You cannot change the permissions associated with the `System` account.
- Do not change the permissions for the `HaUser` account.
- The User Management permission does not allow you to change your own permissions. To change your own permissions, log in to Tiered Storage Manager with the `System` account or as another user with the User Management permission.

To set permissions for a user:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**.
2. In the navigation area, expand the **Users** object tree, and then select the ID of the user whose permissions you want to set.

The *User-ID* subwindow appears in the information area.

3. Click the **Change Permission** button.

   The Change Permission - *User-ID* dialog box appears.

4. For each application, select the permissions you want to set for the user, and then click the **OK** button.

Viewing or editing a user profile

With the User Management permission, you can view and edit the profile of any user. Without the User Management permission, you can view and edit your own profile only. Note that user IDs cannot be changed.

You cannot edit the profiles of users who log in using an external authorization server because their accounts have not been registered in the Hitachi Storage Command Suite products. You must use the Windows Active Directory to edit the profile of these users.

Viewing or editing another user's profile

To view or edit another user's profile:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**.
2. In the navigation area, expand the **Users** object tree, and then select the ID of the user whose profile you want to view or edit.

   The *User-ID* subwindow appears in the information area.

3. Click the **Edit Profile** button.

   The Edit Profile - *User-ID* dialog box appears.

4. Edit **Full Name**, **E-mail**, and **Description** as necessary, and then click the **OK** button.

Viewing or editing your own profile

To view or edit your own profile:

1. In the **Explorer** menu, choose **Settings** and then **User Profile**.

   The User Profile subwindow appears in the information area.

2. To edit the profile, click the **Edit Profile** button.

   The Edit Profile - *User-ID* dialog box appears.

3. Edit **Full Name**, **E-mail**, and **Description** as necessary, and then click the **OK** button.

Changing passwords

With the User Management permission, you can change the password of any user (except for users who are externally authenticated). Without the User Management permission, you can change your own password only. You cannot change the passwords of users for whom external authentication is enabled because an external authentication server manages such passwords.

When changing the password for the `System` account when the Tiered Storage Manager server is running in a cluster configuration, set the same password for all nodes in the cluster.

Caution: If Hitachi Storage Command Suite products and Storage Navigator Modular 2 are installed on the same system, Common Component centrally manages the user accounts for both. Therefore, if you change the password for a user ID in either a Hitachi Storage Command Suite product or in Storage Navigator Modular 2, the password for the other is also changed.

To prevent unwanted access, we recommend that you change this password regularly.

Changing another user's password

To change another user's password:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**.
2. In the navigation area, expand the **Users** object tree, and then select the user ID whose password you want to change.

   The *User-ID* subwindow appears in the information area.
3. Click the **Change Password** button.

   The Change Password - *User-ID* dialog box appears.
4. Type the new password in the **New Password** and **Verify Password** text boxes, and then click the **OK** button.

   Passwords must satisfy all set password conditions, such as the required character combinations or minimum number of characters. For details about password conditions, see Setting password conditions.

Changing your own password

To change your own password:

1. In the **Explorer** menu, choose **Settings** and then **User Profile**.

   The User Profile subwindow appears in the information area.
2. Click the **Change Password** button.

   The Change Password - *User-ID* dialog box appears.
3. Type your current password in the **Old Password** text box. Type your new password in the **New Password** and **Verify Password** text boxes, and then click the **OK** button.

   Passwords must satisfy any set password conditions, such as the required character combinations or minimum number of characters. For details about password conditions, see Setting Password Conditions.

Setting password conditions

To prevent third parties from guessing passwords, you can set a password policy (including the minimum character count and character combinations). The password policy is applied when a user is added or when a password is changed.

When the password policy is changed using the Web client, the `security.conf` file is overwritten. If you have manually edited the `security.conf` file, we recommend that you back it up before changing the password policy.

If Tiered Storage Manager server is used in a cluster environment, set the password policy on the server.

For details about how to set the `security.conf` file and server-side password policy, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*

When using an external authentication server to authenticate users, the settings on the external authentication server are used to determine the combination of character types that can be used for passwords.

To set password conditions:

1. In the **Explorer** menu, choose **Administration**, and then **Security**. Then, in the navigation area, choose **Password**.

   The Password subwindow, which displays the set password conditions, appears in the information area.

2. Click the **Edit Settings** button.

   The Password dialog box appears.

3. Specify the password conditions, and then click the **OK** button.

Setting the automatic locking count for user accounts

To prevent unwanted access, a user's account can be automatically locked if an invalid password is repeatedly specified. Locking can be set to occur automatically when login fails a given number of times.

In this option, a login error on other Hitachi Storage Command Suite products using the single sign-on function is also counted as an error. For example, if the error count is set to 3 and the user fails to log in with Tiered Storage Manager, Device Manager, and Tuning Manager once each, the user account is automatically locked.

A locked user account cannot be used to log in to any Hitachi Storage Command Suite products until it is unlocked.

When the automatic locking is changed using the Web client, the `security.conf` file is overwritten. If you have manually edited the `security.conf` file, we recommend that you back it up before changing this setting.

If Tiered Storage Manager server is used in a cluster environment, set the automatic locking on the server.

For user accounts that log in by linking to an external authorization server, the settings on the external authentication server are used to control automatic locking.

By default, the `System` account is not automatically locked. To automatically lock the `System` account, you must edit the `user.conf` file on the Tiered Storage Manager server.

For details about how to set the `security.conf` file, `user.conf` file, and server-side automatic locking, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

If you want to lock an account that is managed by an external authentication server, you have to manually lock it. For details about how to manually lock an account, see Changing the lock status of a user.

To set the automatic locking of a user account:

1. In the **Explorer** menu, choose **Administration**, and then **Security**. Then, in the navigation area, choose **Account Lock**.

   The Account Lock subwindow, which displays the currently set error count, appears in the information area.

2. Click the **Edit Settings** button.

   The Account Lock dialog box appears.

3. In the Account Lock dialog box, change the settings for the automatic locking of user accounts, and then click the **OK** button.

Changing the lock status of a user

User accounts that were automatically locked due to failed logins can be unlocked, and the accounts of illicit users can be locked. In addition, note the following:

- By default, the `System` account cannot be locked. To change the lock status of the `System` account, you must edit the `user.conf` file on the Tiered Storage Manager server. For details about the `user.conf` file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.
- For users for which external authentication is enabled, the settings on the external authentication server take effect.
- Users cannot lock their own user account, even if they have the User Management permission.
- If you lock the user account of a user who is currently logged in, the user will not be able to perform any operations. Before locking a user account, make sure the user is not logged in to any Hitachi Storage Command Suite products.
- You cannot manually lock the account of a user who has User Management permission (including the `System` account) if the accounts of all other users with User Management permission are already locked.

- If all user accounts are locked and no one can log in, unlock an account by executing the `hcmdsunlockaccount` command on the Device Manager server. For details about this command, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.
- In Hitachi Storage Command Suite products, you cannot change the lock status of a user that logs in by linking to an external authorization server because those user accounts have not been registered in the Hitachi Storage Command Suite products.
- When you change the user authentication from external authentication to user authentication common to Hitachi Storage Command Suite products and the user account password has not been set, the user account will be locked. If you set a password, the user account will be unlocked.

To change the lock status of a user account:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**. In the navigation area, select **Users**.

   The Users subwindow appears.
2. In **User List**, select the user whose account status you want to change. To change the lock status of all user accounts, select the check box in the title bar. If you select a user account that is locked because no password has been set, the user account cannot be unlocked.

   To change the lock status of all user accounts, select the check box in the title bar.

   A window appears for confirming that you want to change the lock status for the specified user account.
3. Click the **Lock Users** or **Unlock Users** button.

   A dialog box appears.
4. Confirm that you want to change the lock status and click **OK**.

Deleting a user

User accounts that are no longer necessary can be deleted.

When deleting user accounts, keep the following conditions in mind:

- You cannot delete the `System` account.
- The `HaUser` account must not be deleted. If you delete this account, the Device Manager agent cannot transmit information to the Device Manager server.
- If you delete a user who is currently logged in, that user will be unable to continue operations. Before deleting any account, make sure the user is not logged in to any Hitachi Storage Command Suite Common products.

To delete a user from the Hitachi Tiered Storage Manager server:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**. In the navigation area, select **Users**.

   The Users subwindow appears.

2. In **User List**, select the check box of each user that you want to delete, and then click the **Delete Users** button.

   The Delete User - Confirmation dialog box appears.

Note: To delete a user after viewing the user's profile, select the desired user in **User List** to open the *User-ID* subwindow. In the *User-ID* subwindow, you can click the **Delete User** button to delete that account.

3. Click the **OK** button to delete the account, or the **Cancel** button to cancel deletion.

Managing user authentication methods

To manage users accounts for Hitachi Storage Command Suite products together with other application programs, you can use an external authentication server (an LDAP directory server, a RADIUS server, or a Kerberos server for user authentication) to authenticate user accounts. To do this, create user accounts for the Hitachi Storage Command Suite products that are the same as the user accounts managed by the external authentication server, and then use the external authentication server to authenticate the user IDs and passwords. You can specify whether each user account is authenticated by the Hitachi Storage Command Suite products or by the external authentication server. The `System` account and `HaUser` accounts cannot be authenticated by an external authentication server.

To enable external authentication for users, you need to set up the server and register user accounts. To do this, perform the procedure in Linking users to an external authentication server, and then perform the procedure in Changing the authentication method of a user. To disable external authentication, specify that Hitachi Storage Command Suite products be used for authentication by performing the procedure in Changing the authentication method of a user.

If you are linking to an external authentication server, and you are also linking to an external authorization server, you can manage user accounts for Hitachi Storage Command Suite products in by authorization group. For more information, see Enabling external authentication using an external authorization server.

Linking users to an external authentication server

The following procedure describes how to link users to an external authentication server. If the users for Hitachi Storage Command Suite products are already registered, perform only steps 1 and 2.

To link users to an external authentication server:

1. On the Tiered Storage Manager management server, specify the settings for linking to an external authentication server.

   For details about requirements for an external authentication server and how to specify the settings for linking to the server, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

2. In the external authentication server, register the IDs and passwords of users for Hitachi Storage Command Suite products.

   The specified user IDs must consist of character strings that can be used in Hitachi Storage Command Suite products. For details about specifiable characters, see the online Help.

   To register a password in a Hitachi Storage Command Suite product, the password must consist of a character string that can be used in Hitachi Storage Command Suite products.

3. Create Hitachi Storage Command Suite products users by using the user IDs registered in step 2.

   Specify the same user IDs registered in the external authentication server as the user IDs common to the Hitachi Storage Command Suite products. If you use a RADIUS server or a Kerberos server that is in a realm other than the realm specified as the default in step 1, specify a character string that contains the realm for the user ID. For example, if the user name that is authenticated by a RADIUS server or Kerberos server is user1 and the realm is example.com, use @ as the separator and specify user1@example.com as the user ID. If multiple user IDs have the same user name and different realms, each user ID is authenticated as a separate user ID.

   Setting passwords is optional. Users authenticated by the external authentication method are authenticated by the passwords registered in the external authentication server. Note that you can set passwords in advance because, if passwords have not been set, the user account will be locked when you change the authentication method from external authentication to user authentication common to Hitachi Storage Command Suite products.

4. Specify permissions for the users registered in step 3.

   For details about how to specify user permissions, see Setting permissions for a user.

Changing the authentication method of a user

To change the method used to authenticate a user:

1. From the **Explorer** menu, choose **Administration** and then **Users and Permissions**.
2. In the navigation area, expand the object tree, and then select **Users**.

   The Users subwindow appears.

3. In the list of users, select the check box of each user for whom you want to enable or disable the ability to link to the external authentication server, and then click the **Change Auth** button.

   The Change Authentication Method dialog box appears.

4. Specify the authentication method, and then update the information.

   The **Authentication** column displayed for each user in the Users subwindow is refreshed.

When you change the user authentication from external authentication to user authentication common to Hitachi Storage Command Suite products and the user account password has not been set, the user account will be locked. To unlock the account, you must set a password.

Enabling external authentication using an external authorization server

If an external authorization server is linked to, you can set permissions for Hitachi Storage Command Suite products in the Windows Active Directory group (authorization group). You do not need to register individual user accounts in Hitachi Storage Command Suite products. However, if you register a user account in a Hitachi Storage Command Suite product, the permission assigned to the user account, rather than that of the authorization group to which that user belongs, is applied even if linking to an external authorization server is enabled. When a group is registered in an authorization group, the permissions set for upper-level groups are not inherited. Only the permissions of the authorization groups registered in Hitachi Storage Command Suite products will be applied.

User IDs to be authenticated per authorization group must consist of a character string that can be used in Hitachi Storage Command Suite products.

To link to an external authorization server:

1. On the Device Manager server, specify settings for linking to an external authorization server.

   For details about the requirements for external authorization servers and settings for enabling linking to an external authorization server, see the *Hitachi Device Manager Software Server Configuration and Operation Guide*.
2. In the **Add Groups - *domain-name*** dialog box, register an authorization group that will use the Hitachi Storage Command Suite products.
3. In the **Change Permission - *domain-name*** dialog box, specify permissions for the authorization group registered in step 2.

Registering an authorization group

To enable linking with an external authorization server, you must register all authorization groups in the Hitachi Storage Command Suite products. You must not only register authorization groups that are already registered in the authorization server, but you must also register any other authorization groups by first registering them in the Hitachi Storage Command Suite products and then registering them in the authorization server.

By default, *All Resources* is assigned to the authorization group as the resource group.

To register an authorization group:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**.

2. In the navigation area, expand the **Groups** object tree and select *domain-name*.

   The *domain-name* subwindow appears.

3. Click the **Add Groups** button.

   The Add Groups - *domain-name* dialog box appears.

4. In the **Distinguished Name** text field, enter the distinguished name of the authorization group.

   A maximum of 20 distinguished names can be registered at a time. Distinguished names must be between 1 - 250 characters. Names are not case-sensitive. The appearance of *RDN* at the beginning of a distinguished name indicates an authentication group name.

5. Click the **Check DN** button to make sure the distinguished name is registered on the external authentication server.

   Verifying registration by using the Check DN button is optional. If a distinguished name is not registered on an external authentication server, a message appears. However, you will still be able to register the authorization group in Hitachi Storage Command Suite products.

6. Click the **OK** button.

   The dialog box closes and the authorization group is registered.

Setting permissions for an authentication group

After registering an authentication group in a Hitachi Storage Command Suite product, assign permissions to the group. For users who belong to multiple authorization groups, all the permissions assigned to the groups are applied to the user.

To set permissions for a registered authorization group:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**.
2. In the navigation area, expand the **Groups** object tree and select *domain-name*.

   The *domain-name* subwindow appears.

3. Click the **Change Permission** button.

   The Change Permission - *group-name* dialog box appears.

4. Select the permissions to be specified for the group for each application and then click **OK** to save the settings.

Deleting an authorization group

Users with user management permissions can delete authorization groups registered in Hitachi Storage Command Suite product.

To delete a registered authorization group:

1. In the **Explorer** menu, choose **Administration** and then **Users and Permissions**.

2. In the navigation area, expand the **Groups** object tree and select *domain-name*.

   The *domain-name* subwindow appears.

3. In the **Group List**, select the check box of each group you want to delete and then click **Delete Groups**.

   The Delete Groups - *domain-name* subwindow appears.

   To select all groups, select the check box for all groups in the title bar.

   Important: If you view information about a group before you delete it, select *domain-name* in the **Groups** object tree and then *group-name* to open the group-name subwindow. If you click the Delete Groups button in this window, you can delete the displayed authorization group.

4. To delete the group, click **OK**.

   To cancel the deletion, click the **Cancel** button.

Setting a message for the warning banner

Users with User Management permission can set a message that is displayed in the login window as a *warning banner*. You can use a limited set of HTML tags to format the warning banner. For details about the usable HTML tags, see the online Help.

Note: The warning banner message can also be set by using the `hcmdsbanner` command. When the `hcmdsbanner` command is used to set the message, there are no limitations on the usable HTML tags. Additionally, a different message can be set for each locale. For details about using the `hcmdsbanner` command to set the warning banner, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

If a message set using the `hcmdsbanner` command contains HTML tags that cannot be used in the Web client, the message cannot be edited in the Web client.

To edit the warning banner message:

1. In the **Explorer** menu, choose **Administration**, and then **Security**. Then, in the navigation area, choose **Warning Banner**.

   In the information area, the Warning Banner subwindow appears. The Warning Banner subwindow displays the message that has been set.

2. Click the **Edit Message** button.

   The Edit Message dialog box appears.

3. In the **Message** text box, edit the message. To delete the existing message, click the **Delete** button.

4. Click the **Preview** button.

   If the message has been specified correctly, the message is displayed as formatted HTML in the **Preview** area. If an unsupported tag is used or there is an HTML syntax error, an error message is displayed and the **Preview** area is blank. If you want to edit the message again, click the **Preview** button once more.

5. To accept the message, click the **OK** button.

Setting up storage domains

If volumes that you want to use with Tiered Storage Manager are in a storage subsystem registered in Device Manager, you can register that storage subsystem in Tiered Storage Manager.

Creating a storage domain

To register a storage subsystem registered in Device Manager to Tiered Storage Manager, you must create a new storage domain.

To create a new storage domain:

1. In the **Explorer** menu, choose **Resources** and then **Domains**.

   The Storage Domains subwindow appears in the application area.

2. Click the **Create Domain** button.

   The Create Domain dialog box appears.

3. Enter **Name** and **Description**, and select a storage subsystem from **Domain Ctrl**.

4. Click the **OK** button.

   The Create Domain dialog box is displayed, and the storage domain is refreshed. Clicking the **X** button in the upper right corner of the window does not cancel the processing. When the processing is completed, a processing result dialog box appears.

5. Click the **Close** button.

   The processing result dialog box closes, and the information in the navigation area and application area is updated.

Checking a storage domain's properties

To check information such as the migration groups or storage tiers registered for a storage domain, display the *storage-domain-name* subwindow.

To display the *storage-domain-name* subwindow:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area or Storage Domains subwindow, choose the desired *storage-domain-name* link.

Editing storage domain properties

The Storage Domains subwindow or the *storage-domain-name* subwindow can be used to change the properties of a created storage domain.

To use the Storage Domains subwindow to change storage domain properties:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**.

   Then in the Storage Domains subwindow, click the  icon for the desired storage domain.

   The Edit Domain dialog box appears.

2. Edit **Name** and **Description** as necessary, and then click the **OK** button.

To use the *storage-domain-name* subwindow to change storage domain properties:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area or Storage Domains subwindow, choose the desired *storage-domain-name* link.

   The *storage-domain-name* subwindow appears.

2. Click the **Edit Domain** button.

   The Edit Domain dialog box appears.

3. Edit **Name** and **Description** as necessary, and then click the **OK** button.

Deleting a storage domain

If you no longer need to manage a storage subsystem by using Tiered Storage Manager, you must delete the storage domains in the storage subsystem. When a storage domain is deleted, the volumes within it are not deleted. Note that storage domains cannot be deleted if they contain unfinished tasks.

The Storage Domains subwindow or *storage-domain-name* subwindow can be used to delete a storage domain.

To use the Storage Domains subwindow to delete a storage domain:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**.

   Then in the Storage Domains subwindow, click the  icon for the desired storage domain.

   The Delete Domain dialog box is displayed.

2. Select the check box **Yes. I have read the above, and still want to continue**, and then click the **OK** button.

   A message is displayed when the storage domain has been deleted.

3. Click the **Close** button.

To use the *storage-domain-name* subwindow to delete a storage domain:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then in the navigation area or Storage Domains subwindow, choose the desired *storage-domain-name* link.

   The *storage-domain-name* subwindow appears.

2. Click the **Delete Domain** button.

   The Delete Domain dialog box is displayed.

3. Select the check box **Yes. I have read the above, and still want to continue**, and then click the **OK** button.

   A message is displayed when the storage domain has been deleted.

4. Click the **Close** button.

Refreshing storage domains

When the information about a storage domain is not the most recent information, a warning message appears in the window. If this message is displayed, refresh the storage domain.

To refresh a storage domain:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**.

   The Storage Domains subwindow appears.

2. Click the  icon for the desired storage domain.

   The Refresh Domain dialog box asks you to confirm the refresh operation.

   You can also refresh the information by clicking the **Refresh Domain** button in the *storage-domain-name* subwindow.

3. Click the **OK** button.

   The storage domain is refreshed. When refresh processing is completed, a processing result message appears.

4. Click the **Close** button.

Setting up storage tiers

This section explains how to manage storage tiers, which are migration targets.

When using the following search conditions to create a storage tier, make sure that the external storage subsystem has been registered in Device Manager:

- Subsystem
- Subsystem Vendor
- Subsystem Display Model
- Subsystem Serial Number

The search results that are displayed in Tiered Storage Manager are different, depending on whether all of the external storage subsystems are registered in Device Manager. Because the naming conventions for the RAID information and the Device Manager information are different, if the above search conditions are used to create a storage tier in an unregistered subsystem, and then that subsystem is registered in Device Manager, the search results will no longer display the volumes contained in the storage tier mentioned above.

Creating a storage tier by using the Create Tier dialog box

This section explains how to create a group of volumes or pools by using the Create Tier dialog box.

To create a group of volumes or pools by using the Create Tier dialog box:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. In the navigation area, choose the desired *storage-domain-name* link, and then the **Storage Tiers** page.

   The Storage Tiers subwindow appears.

2. Click the **Create Tier** button.

   The Create Tier dialog box appears.

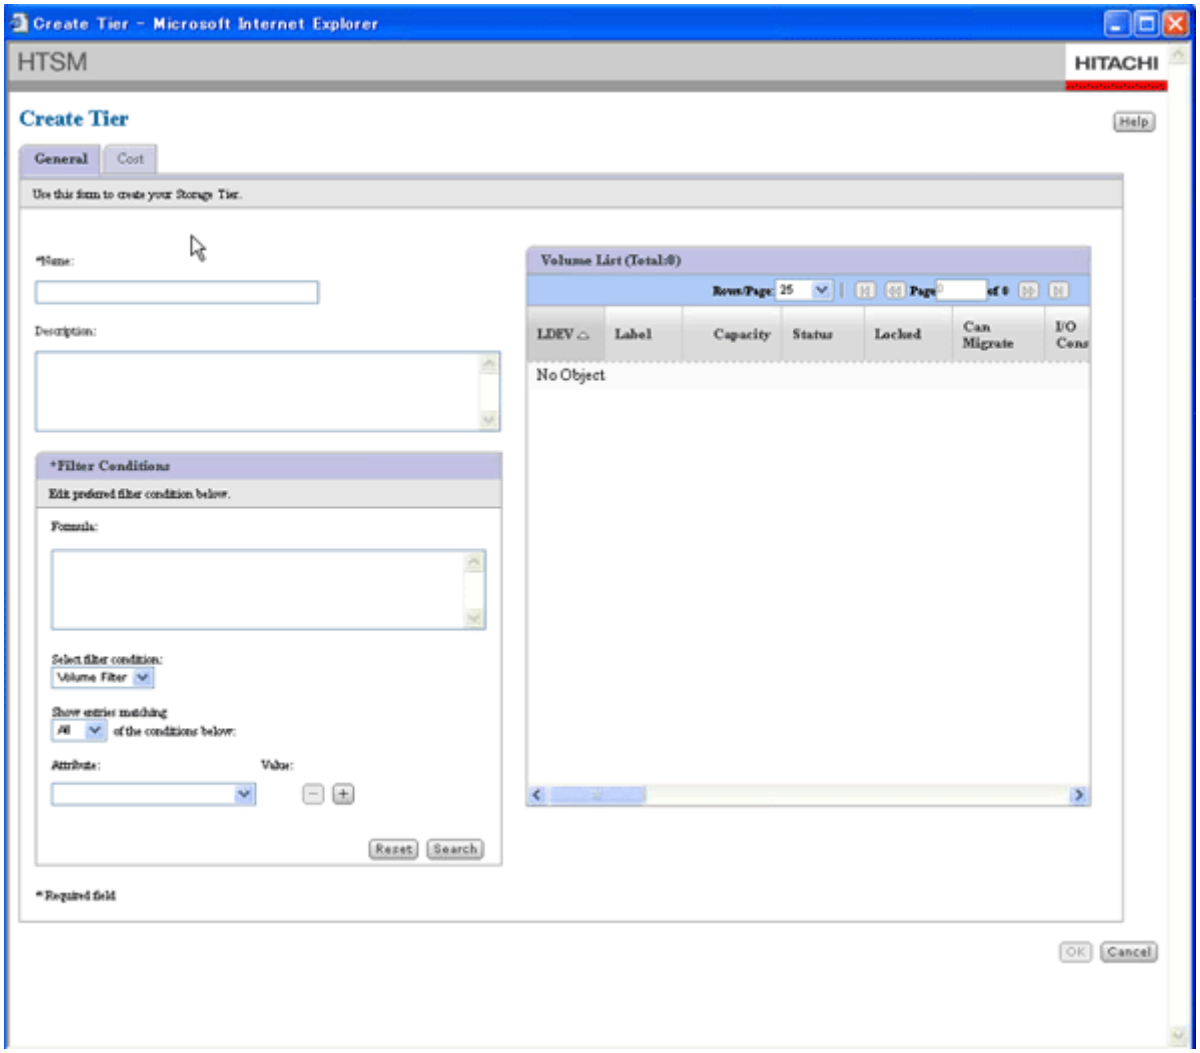

Figure 3-2: Create Tier - General tab dialog box

3. With the **General** tab selected, specify **Name** and **Description**.
4. In the **Select filter conditions** drop-down list in the **Filter Conditions** area, select to search for either volumes or pools.

   Volume filter conditions and pool filter conditions cannot be specified at the same time.

5. Specify attributes, operators, and values for the volume or pool filter conditions.
6. Click the **Search** button.

   The search results are displayed in **Volume List** or **Pool List**.

7. Click the **OK** button to create the tier.

   The Create Tier - Creating dialog box appears, followed by the processing result dialog box.

8. Click the **Close** button.

Creating a storage tier by specifying volume filter conditions

This section explains how to specify volume filter conditions in the Search subwindow and then create a storage tier. You cannot specify pool filter conditions in this subwindow.

To specify volume filter conditions and then create a storage tier:

1. In the **Explorer** menu, choose **Search**.

   The navigation area displays items for constructing a search condition.

2. Select the relevant items from the pull-down menus.
3. Click the **Search** button.

   The application area displays the Search subwindow containing the search results.

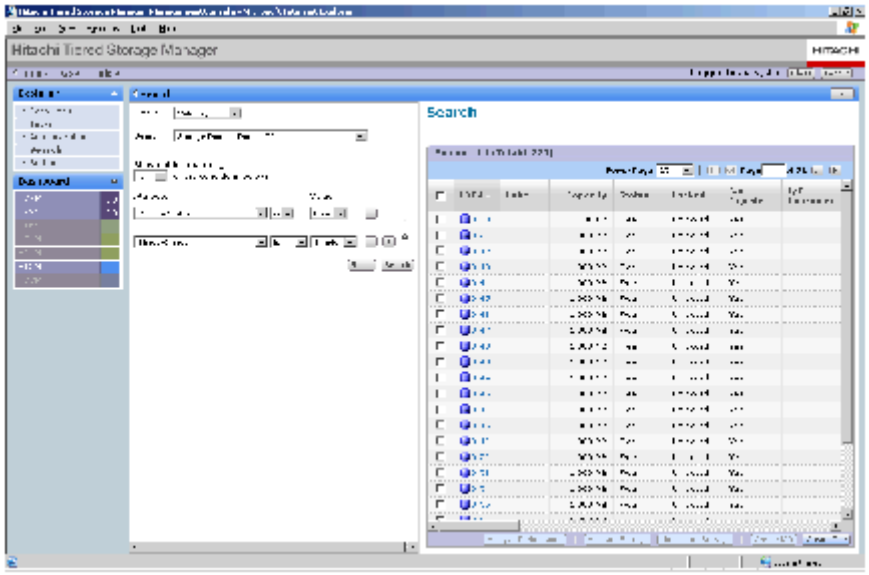

Figure 3-3: Search subwindow

If you want to change the search conditions, edit them, and then click the **Search** button again.

Assigning or releasing a path for a volume

In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired storage-domain-name, **Storage Tiers**, and then the desired storage-tier-name link.

Select the check box for the desired volume, and then click the **Allocate Storage** button (or **Unallocate Storage** button).

Checking volume performance information

To check the performance information of a volume, select the check box for the volume, and then click the **Analyze Performance** button. The Tuning Manager Historical Report dialog box appears.

4. Click the **Create Tier** button.

   The Create Tier dialog box appears.

5. In the Search subwindow, set up the necessary information.

   You can change the search conditions in this step.

6. Click the **OK** button.

   The Create Tier - Creating dialog box appears, followed by the processing result dialog box.

7. Click the **Close** button.

Assigning cost attributes to a storage tier

You can assign cost attributes to a storage tier per unit volume or per DP-volume.

When you move from the **General** tab to the **Cost** tab, the application validates the existing entries in all options and fields, including the search criteria. If there is an error in any option or field, the dialog box does not switch over to the target tab pane.

To assign cost attributes:

1. In the **Create Tier** pane, select the **Cost** tab.

   If your tier search properties validate, the **Cost** dialog box appears.

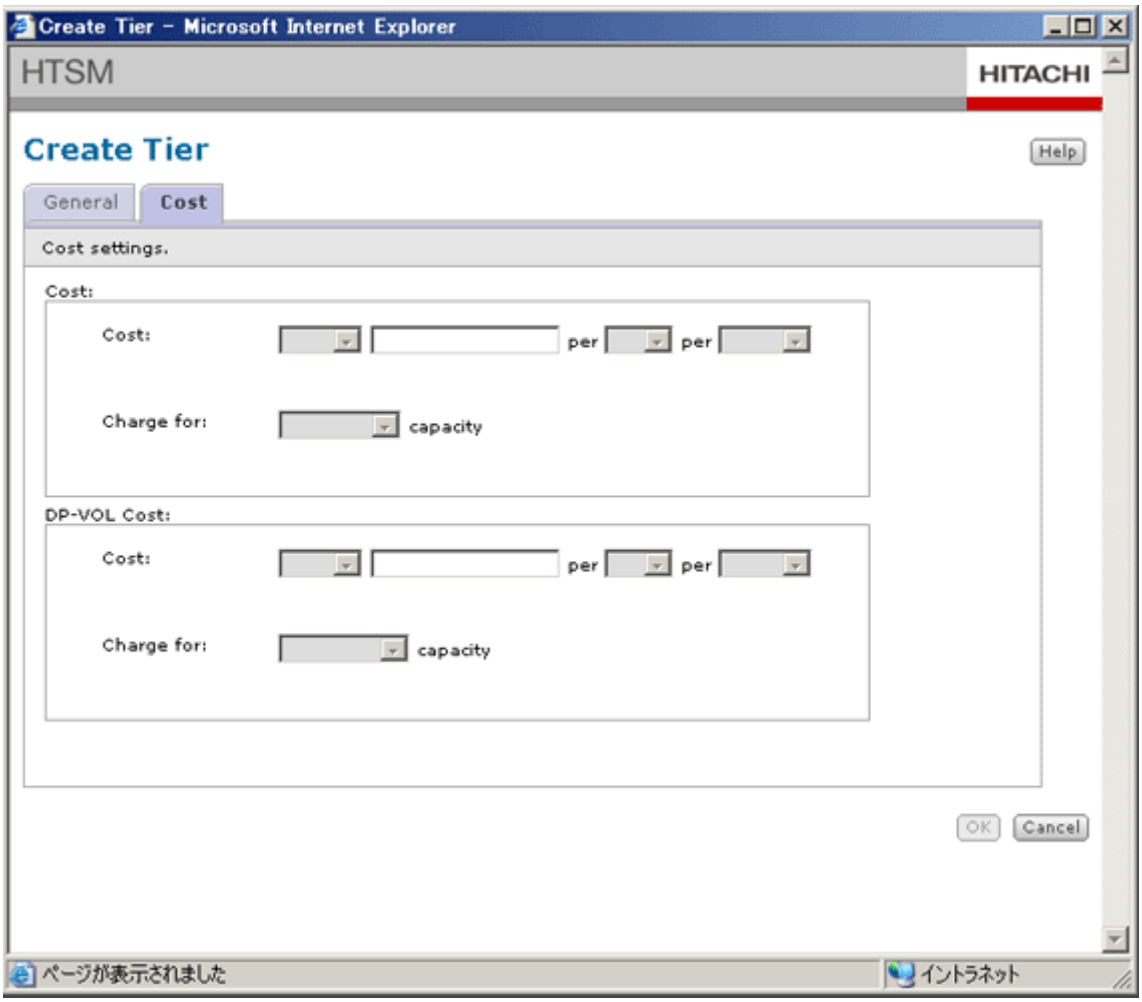

Figure 3-4: Create Tier - Cost tab dialog box

2. Using the drop-down lists, assign values for cost per unit.

Checking a storage tier's properties

Storage tier properties can be checked from the *storage-tier-name* subwindow.

To check a storage tier's properties:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link, **Storage Tiers**, and then the desired *storage-tier-name* link.

   The *storage-tier-name* subwindow appears. The *storage-tier-name* subwindow can also be displayed by clicking the *storage-tier-name* link in the Storage Tiers subwindow or the *storage-tier-name* link in the **Storage Tiers** page of the *storage-domain-name* subwindow.

Editing a storage tier's properties

To change the name, description, or filter conditions (attribute information) of an existing storage tier, change the storage tier properties in the Edit Tier dialog box.

To edit the properties of a storage tier:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link, and then **Storage Tiers**.

   The Storage Tiers subwindow appears.

2. Click the  icon for the desired storage tier.

   The Edit Tier - *storage-tier-name* dialog box appears. The Edit Tier dialog box can also be displayed by clicking the **Edit Tier** button in the *storage-tier-name* subwindow.

3. Edit the storage tier properties.

   The properties that can be edited include **Name**, **Description**, **Cost**, and **Filter Conditions**.

4. Click the **OK** button.

   The Edit Tier result dialog box appears.

5. Click the **Close** button.

Assigning or releasing a path for a volume in a storage tier

From the *storage-tier-name* subwindow, you can link to Device Manager Web Client to assign or release a path for a volume in a storage tier. To link to Device Manager Web Client, you must have the Modify permission for both Tiered Storage Manager and Device Manager.

To assign a path:

1. In the **Explorer** menu, select **Resources** and then **Storage Domains**.
2. In the navigation area, select **Storage Domain Name**, **Storage Tiers**, and then **Storage Tier Name** links.
3. In the **Volume List** page, select the check box for the volume to which you want to assign a path.
4. Click the **Allocate Storage** button.

   The Add Device Manager Add Storage - Step 1 of 4 window appears as a separate window.

5. Follow the on-screen instructions to assign a path to the volume.

To release a path:

1. In the **Volume List** page, select the check box for the volume whose path you want to release.
2. Click the **Unallocate Storage** button.

   The Cancel Device Manager Unallocate Storage - Confirmation window appears as a separate window. For details about using Device Manager, see the *Device Manager online Help*.

Deleting a storage tier

To delete a storage tier:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link, **Storage Tiers**, and then the desired *storage-tier-name* link.

   The *storage-tier-name* subwindow appears.

2. Click the **Delete Tier** button.

   The Delete Tiers dialog box appears.

3. In the Delete Tiers dialog box, select the check box **Yes. I have read the above, and still want to continue.**

4. Click the **OK** button.

   A dialog box appears, indicating that the storage tier is being deleted. When deletion processing is completed, the Delete Tiers result dialog box appears automatically.

5. Click the **Close** button.

Creating volumes

This section explains how to search for free capacity in storage subsystems and then create volumes based on the search results.

Searching for free space and creating volumes

To create volumes, you must create a Volume Creation Task.

To create a Volume Creation Task:

1. In the **Explorer** menu, choose **Search**.

   The navigation area displays items for constructing filter conditions.

2. Select **Free Space(s)** for **Find all**.

3. Select the relevant items from the pull-down menus.

4. Click the **Search** button.

   The application area displays the Search subwindow, which contains the search results.

5. Click the radio button for the free space to be used for the volumes to be created.

6. Click the **Create Volume** button.

   The Create Volume dialog box appears.

7. Specify **Emulation Type**, the volume size, the number of volumes, and the required task options, and then click the **Confirm** button.

   If you specify a value other than OPEN-V for the emulation type, the volumes will be created in the free space of any of the selected array groups.

   If the value `1` after **No. of Volumes** is grayed out, you cannot set the number of volumes. In this case, only one volume will be created.

Click the **Confirm** button to create the Volume Creation Task. The task execution status can be checked in the Tasks subwindow.

If you are using Lightning 9900, Lightning 9900V, Hitachi USP, or Universal Storage Platform V/VM, the capacity of the created volume will be larger than the specified value, because a control area is created on the storage subsystem side when the volume is created. Therefore, you cannot create a volume that is the same size as the free space.

When you create a volume in an SMI-S Enabled subsystem, the volume might be created at the minimum capacity if the specified capacity is smaller than the minimum capacity allowed in the corresponding storage subsystem. This is within the SMI-S provider specifications.

Adding and refreshing storage subsystems

This section explains how to add and refresh storage subsystems.

Registering storage subsystems in Device Manager

To register a storage subsystem in Device Manager:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**.

   The Storage Domains subwindow appears.
2. Click the **Discover Subsystem** button.

   The Add Storage dialog box of Device Manager appears. For details about using Device Manager, see the *Device Manager online Help*.

Refreshing storage subsystems

To refresh a storage subsystem registered in Device Manager:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**.

   The Storage Domains subwindow appears.
2. Click the **Refresh Subsystem** button.

   The Refresh Subsystem dialog box appears.
3. In the drop-down list, select the storage subsystem you want to refresh.
4. Click the **OK** button.
5. Click the **OK** button.

   The storage subsystem registered in Device Manager is refreshed.

Connecting external storage subsystems to a storage domain

This section explains how to connect a storage subsystem registered in Device Manager to Universal Storage Platform V/VM.

Setting up external mapping

To set up external mapping, select an external storage subsystem and then connect it to Universal Storage Platform V/VM. To connect the external storage subsystem, you must create an External Mapping Task.

To create the task:

1. In the **Explorer** menu, choose **Search**.

   The navigation area displays items for constructing filter conditions.
2. Select **Volume(s)** for **Find all**.
3. Select an external storage subsystem for **Search from**.
4. Select the relevant items from the pull-down menus.
5. Click **Search**.

   The application area displays the Search subwindow, which contains the search results.
6. Select a volume to be virtualized, and then click the **Virtualize Storage** button.

   The Virtualize Storage dialog box (default display) appears.
7. In the drop-down list, select a storage subsystem that is already registered as a storage domain.
8. Select the path group. Select **New** to create a new path group.
9. Click the **Refresh Path to External Volume** button.

   A different Virtualize Storage dialog box appears.
10. Select the check box for the desired mapping path.

    In **External Port/WWN**, which is displayed as path group information, port pairs that are not physically connected might also be displayed. Confirm which paths port pairs are physically connected for, and then select the appropriate paths.

    To set up an HSD or LUN:

    The  icon appears next to a path for which an HSD or LUN has not been set up. Click this icon to display the Edit Path dialog box. In this dialog box, you can set up HSDs and LUNs.
11. Click the **Define Path Priority** button.

    The Define Path Priority dialog box appears.
12. Use the or button to set the path priority, and then click the **OK** button.

    The Define Path Priority dialog box closes and the Virtualize Storage dialog box from step 9 appears again.
13. Specify the required task options, and then click the **Confirm** button.

    The External Mapping Task is created. The task execution status can be checked in the Tasks subwindow.
14. Repeat steps 6 through 13 to virtualize additional volumes.

After the External Mapping tasks have been completed successfully, open the task information window and click on the **Component Status** tab. If you plan to use Device Number as a search condition for these volumes in the future, record the virtualized volume's LDEV number. Use the LDEV number as the Device Number search condition parameter.

Setting up migration groups

A set of source volumes for migration is called a *migration group*. Migration groups are also used for executing locking tasks, unlocking tasks, and shredding tasks.

Creating a target volume group by using the Create MG dialog box

This section explains how to create empty migration groups and migration groups that contain volumes.

To create a migration group by using the Create MG dialog box:

1. In the **Explorer** menu, choose **Resources**, and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link, and then the **Migration Groups** page.

   The Migration Groups subwindow appears.

   You can also create a migration group by choosing the **Migration Groups** page in the *storage-domain-name* subwindow.

2. Click the **Create MG** button.

   The Create MG dialog box appears. The Create MG dialog box is composed of three tabbed pages. For details about the items in these pages, click the **Help** button.

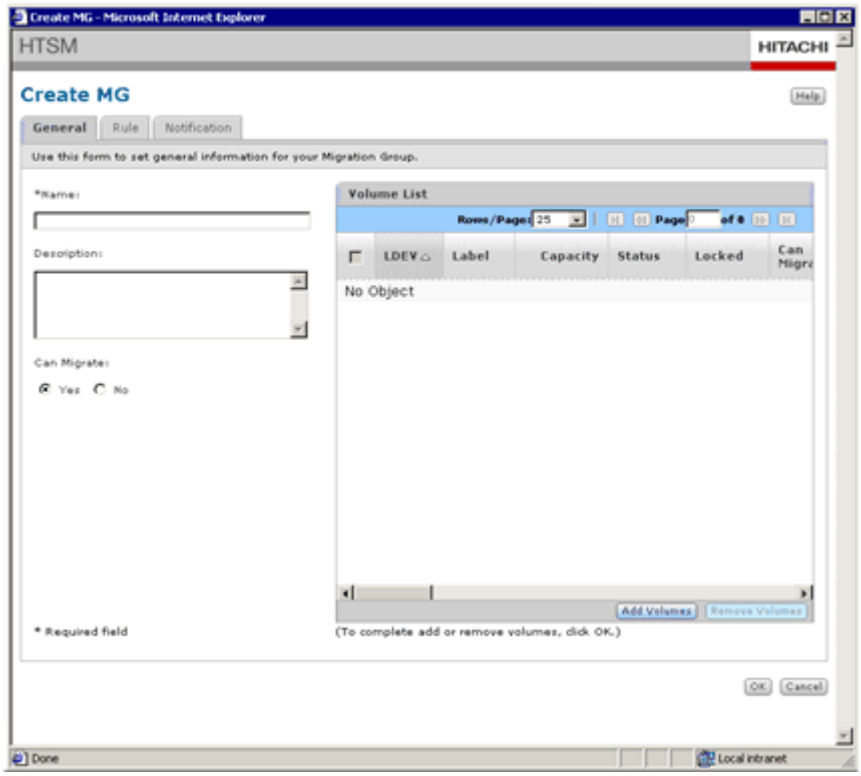

Figure 3-5: Create MG dialog box

3. On each page, set the required items.

   The Create MG dialog box is composed of three tabbed pages. For details about the items in these pages, click the **Help** button.

   To add volumes:

   Clicking the **Add Volumes** button displays the Add Volume dialog box. From the list of volumes, select the check box for each volume to be added, and then click the **OK** button.

   To add to a migration group volume that is already included in another migration group, select the **Allow moving volumes from other Migration Groups** check box. The selected volumes are deleted from the other migration group and added to the created migration group.

4. Click the **OK** button.

   The Create MG - Creating dialog box appears, followed by the processing result dialog box.

5. Click the **Close** button.

Creating a volume group by filtering volumes

This section explains how to create a migration group by using filters to search for source volumes.

To search for volumes and then create a migration group:

1. In the **Explorer** menu, choose **Search**.

The navigation area displays items for constructing a search condition.

2. Select the relevant items from the pull-down menus.
3. Click the **Search** button.

   The application area displays the Search subwindow, which contains the search results.

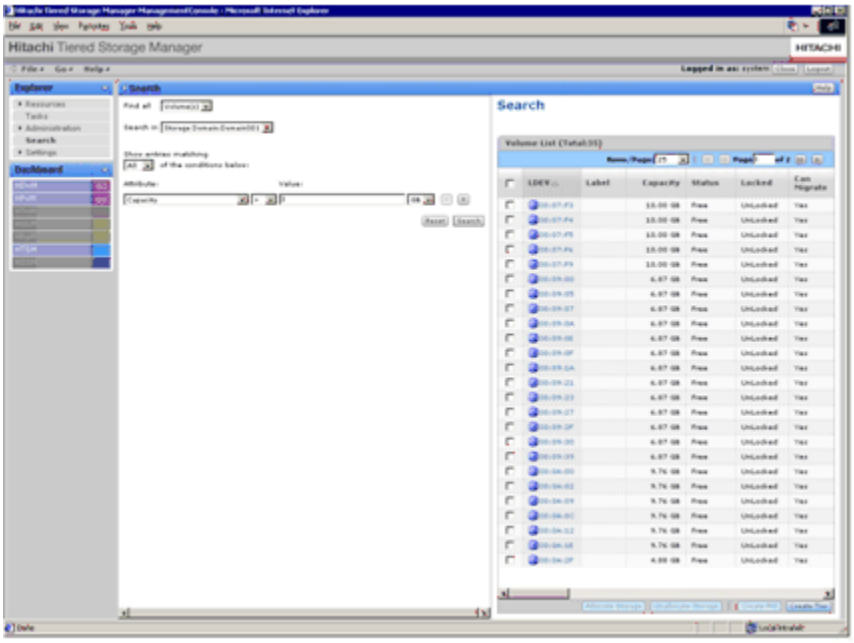

Figure 3-6: Search subwindow

4. Select the volumes to be added to the migration group.

   Select the check box for each desired volume.

   Assigning or releasing paths to volumes

   In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired storage-domain-name, **Storage Tiers**, and then the desired storage-tier-name link.

   Select the check box for the volume and then click the **Allocate Storage** button or **Unallocate Storage** button.

   Checking volume performance information

   To check the performance information of a volume, select the check box for the volume, and then click the **Analyze Performance** button. The Tuning Manager Historical Report dialog box appears.

5. Click the **Create MG** button.

   The Create MG dialog box appears.

6. Specify the information of the migration group to be created.

   Specify the required information in the Create MG dialog box. You can add volumes to the migration group by clicking the **Add Volumes** button.

7. Click the **OK** button.

The Create MG - Creating dialog box appears, followed by the processing result dialog box.

8. Click the **Close** button.

Checking a migration group's properties

You can check the properties of a migration group in the *migration-group-name* subwindow.

To check a migration group's properties:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link, and then **Migration Groups**.

   The Migration Groups subwindow appears.

2. Click the desired *migration-group-name* link.

   The *migration-group-name* subwindow appears.

   The *migration-group-name* subwindow can also be displayed by clicking the *migration-group-name* link in the **Migration Groups** page of the *storage-domain-name* subwindow.

Editing a migration group's properties

To add a volume to an existing migration group, or change a migration group's name, description, LDEV selection rule, or notification email address, change the properties in the Edit MG dialog box.

To change the properties in the Edit MG dialog box:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link, and then **Migration Groups**.

   The Migration Groups subwindow appears.

2. Click the  icon for the desired migration group.

   The Edit MG dialog box appears.

   The Edit MG dialog box can also be displayed by clicking the **Edit MG** button in the *migration-group-name* subwindow.

3. Select the appropriate page and edit the migration group properties as desired.

   For details about the properties that can be edited in each page, click the **Help** button.

4. Click the **OK** button.

   The Edit MG result dialog box appears.

5. Click the **Close** button.

Assigning or releasing a path for a volume in a migration group

From the *migration-group-name* subwindow, you can link to Device Manager Web Client to assign or release a path for a volume in a migration group.

To assign a path:

1. In the **Volume List** page, select the check box for the volume to which you want to assign a path.
2. Click the **Allocate Storage** button.

   The Add Device Manager Add Storage - Step 1 of 4 window appears as a separate window.
3. Assign a path to the volume.

To release a path:

1. In the **Volume List** page, select the check box for the volume whose path you want to release.
2. Click the **Unallocate Storage** button.

   The Device Manager Unallocate Storage - Confirmation window appears as a separate window. For details about using Device Manager, see the *Device Manager online Help*.

Deleting a migration group

Migration groups that are no longer needed can be deleted. Note that when a migration group is deleted, its volumes are not deleted.

To delete a migration group:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link, and then **Migration Groups**.

   The Migration Groups subwindow appears.
2. Click the *migration-group-name* link of the migration group to be deleted.

   The *migration-group-name* subwindow appears.
3. Click the **Delete MGs or Delete MG** button.

   The Delete MGs dialog box appears.
4. In the Delete MGs dialog box, select the check box **Yes. I have read the above, and still want to continue.**
5. Click the **OK** button.

   The Delete MGs dialog box appears. When deletion processing is completed, the Delete MGs result dialog box appears automatically.
6. Click the **Close** button.

Performing migration

After migration groups and storage tiers are created, volume migration can be performed.

There are two ways to migrate volumes:

- Using the Migration Wizard
- Using the Migrate MG dialog box

Migrating volumes by using the Migration Wizard

The following explains how to use the Migration Wizard to perform migration. Migration groups and storage tiers can also be created during operation.

To migrate volumes by using the Migration Wizard:

1. From the global tasks bar, choose **Go,** and then **Migration Wizard**.

   Migration Wizard Step 1/6, which provides an overview of the processing, appears in the Migration Wizard dialog box.

2. Click the **Next** button.

   Migration Wizard Step 2/6 appears, enabling you to select a migration group.

   If the storage domain information is not up to date, an error message is displayed. Refresh the storage domain, and then try again.

3. Select a migration group in the **Migration Group** list box, or to create a migration group, click the **Create MG** button.

   To create a new migration group, click the **Create MG** button.

   Click the **Next** button.

4. Migration Wizard Step 3/6 appears, enabling you to select a storage tier.

5. Select a storage tier.

   An **N** appears under **Comp** for storage tiers that do not have enough target volumes. However, if the array group avoidance rule is set for the migration group, even if **Y** is displayed under **Comp**, the number of target volumes might not be enough.

   To create a new storage tier, click the **Create Tier** button.

6. Click the **Next** button.

   Migration Wizard Step 4/6 appears, enabling you to check the migration pairs automatically created by Tiered Storage Manager.

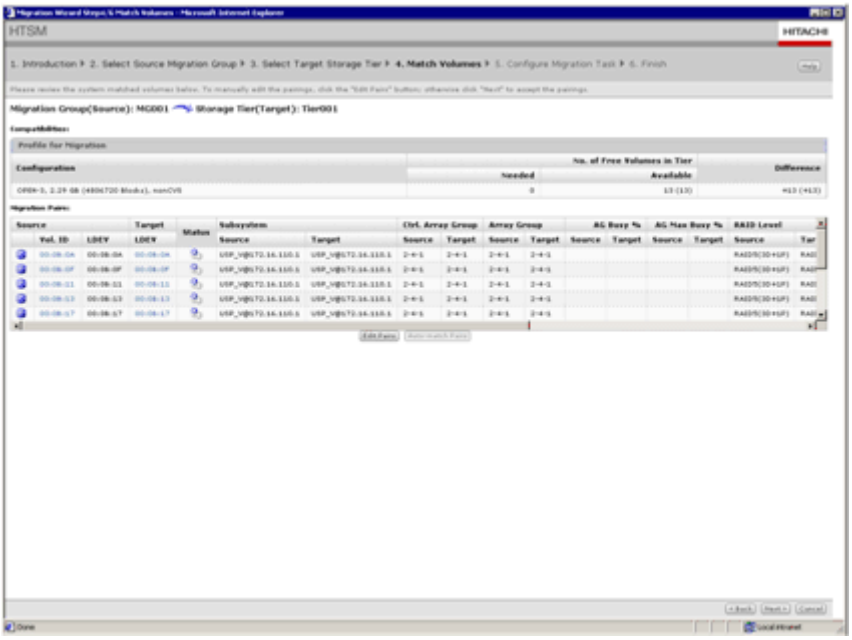

Figure 3-7: Migration Wizard Step 4/6

If a storage tier created from a pool is specified as the migration target, the migration target pool number is displayed instead of the migration target LDEV number.

7. Click the **Auto Match Pairs** button.

   Migration source and target volumes are automatically paired, and then Migration Wizard Step 4/6 appears again.

   To change the migration pairs:

   Click the **Edit Pairs** button. The Edit Migration Pairs dialog box, which enables you to change migration volume pairs, appears.

   To check target migration volume performance information:

   Select the check box for the migration pair, and then click the **Analyze Performance** button. The Tuning Manager Historical Report dialog box appears. You can check the performance information only if the target migration volume is a normal array group.

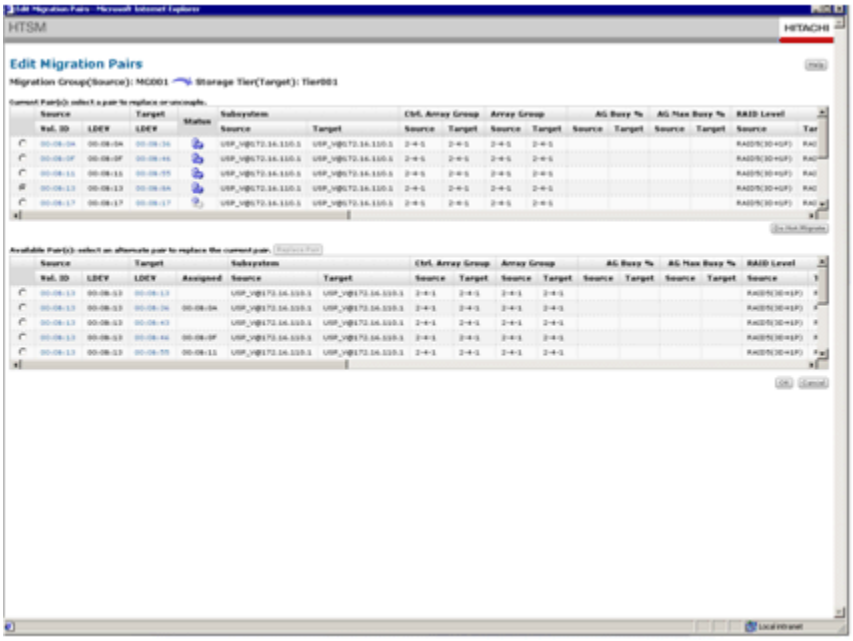

Figure 3-8: Edit Migration Pairs dialog box

8. Select the pair to edit.

   To exclude a migration source volume, select the radio button next to its ID, and then click the **Do Not Migrate** button. If a storage tier created from a pool is specified as the migration target, a check box appears next to the volume ID. If all pairs have the (**Do Not Migrate**) status, however, you cannot create a migration task.

   If you want to create a migration target volume in a storage tier that was created by specifying a free pool capacity for a filter condition, specify the migration target pool. When the migration target is a DP pool, an LDEV number is automatically assigned to the migration target.

   When you finish changing volume pairs, click the **OK** button. The Edit Migration Pairs dialog box closes and Migration Wizard Step 4/6 appears again.

9. Click the **Next** button.

   Migration Wizard Step 5/6 appears, enabling you to set options for the migration task to be created.

   You can set options for executing the created task, erasing data after execution, and, for DP volumes or DP pools, releasing unused space from the target volume (discarding zero data) after migration is completed.

10. Enter the task description and notification email address as necessary, and then select when the created task is to be executed and whether the data at the migration source is to be deleted after the task is executed.

Note: By default, the data on the migration source volumes will not be erased after migration. In order to prevent inappropriate disclosure of information, we recommend that you shred the data on the migration source volumes after migration.

11. Click the **Confirm** button.

    If you set the option for immediate execution in step 9, the created task is executed automatically (migration is executed). When task execution or migration task creation is completed, the end of the Migration Wizard, Step 6/6, is displayed.

12. Click the **Finish** button.

To check the task execution status, choose **Task** in the **Explorer** menu. It is normal for the migration progress indicator to remain at 99% for some time during the final part of the migration.

Migrating volumes by using the Migrate MG dialog box

If the migration groups and storage tiers to be used for migration have already been set up, use the Migrate MG dialog box to perform migration.

To migrate volumes by using the Migrate MG dialog box:

1. In the **Explorer** menu, choose **Resources** and then **Storage Domains**. Then, in the navigation area, choose the desired *storage-domain-name* link.

2. Either click the  icon for the desired migration group or storage tier, or click the **Migrate MG** button.

   To display the  icons, perform one of the following:

   - In the navigation area, choose **Migration Groups**.
   - In the navigation area, choose **Storage Tiers**.
   - In the *storage-domain-name* subwindow, choose the **Migration Groups** page or **Storage Tiers** page.

   To display the **Migrate MG** button:

   - In the navigation area, choose **Migration Groups**, and then the desired *migration-group-name* link.
   - In the navigation area, choose **Storage Tiers**, and then the desired *storage-tier-name* link.

   The Migrate MG dialog box appears.

   If the storage domain information is not up to date, an error message is displayed. Refresh the storage domain, and then try again.

3. Select the migration group from the **Migration Groups (Source)** box, and the storage tier from the **Storage Tiers (Target)** box, and then click the **Check Compatibility** button.

   A check is performed to make sure that the number of migration target volumes is sufficient for the migration source. If there is no problem, the migration source volumes and migration target volumes are automatically paired, and the result is displayed in the **Migration Pairs** table. If there is a problem during the compatibility check, a message will be output to the **Profile for Migration** list, and the **Migration Pairs** table will display a list of the migration pairs that have been set at the migration target.

When a storage tier that was created by specifying a free pool capacity as a filter condition is specified as the migration target storage tier, the migration target pool number is displayed instead of the migration target LDEV number.

To change the migration volume pair:

Click the **Edit Pairs** button. The Edit Migration Pairs dialog box, which enables you to change migration volume pairs, appears.

Checking target migration volume performance information:

Select the check box for the migration pair, and then click the **Analyze Performance** button. The Tuning Manager Historical Report dialog box appears. You can check the performance information only if the target migration volume is a normal array group.

4. Set **Task Description**, **Notification Email**, and other options.

   You can set options for executing the created task, erasing data after execution, and, for DP volumes or DP pools, releasing unused space from the target volume (discarding zero data) after the migration is completed.

5. Click the **OK** button.

   A migration task is created. If you set the option **Execute immediately after confirmation**, the created task is executed automatically (migration is executed). If **Execute immediately after confirmation** was not set, the task is not executed.

   When task execution or migration task creation is completed, the Migrate MG dialog box appears notifying you that the task was completed successfully.

6. Click the **Close** button to close the dialog box.

To check task execution status, choose **Tasks** from the **Explorer** menu. It is normal for the migration progress indicator to remain at 99% for some time during the final part of the migration.

Working with tasks

To execute a task that is on standby, check the execution status of a task, cancel a task, or temporarily stop a task, choose **Tasks** from the **Explorer** menu.

Executing tasks

To execute tasks that are on standby:

1. In the **Explorer** menu, choose **Tasks**.

   The Tasks subwindow appears in the application area.

   The Tasks subwindow displays a list of tasks that have been created in the past 7 days.

2. Select the check box for each task to be executed, and then click the **Execute Tasks** button.

If any of the selected tasks is not on standby, the Execute Tasks dialog box is displayed to confirm task execution. To continue the operation, click the **OK** button.

A dialog box appears indicating that the tasks are being executed. When execution is completed, a dialog box appears.

3. Click the **Close** button.

   The Execute Tasks dialog box closes, and the Tasks subwindow appears again. At this time, the Tasks list will be updated in the Tasks subwindow.

Checking the execution status of tasks

After tasks are executed, some time is required until they complete. The task execution status can be checked from the **Tasks** menu.

To update the task status list to show the latest status, click the **Refresh Tasks** button in the application bar.

To check the execution status of a task:

1. In the **Explorer** menu, choose **Tasks**.

   The Tasks subwindow appears in the application area.

2. Choose the *task-ID* link of the task whose execution status you want to check.

   The *task-ID* dialog box is displayed.

   To check the task status of each volume, choose the **Component Status** page.

Searching for tasks

To search for a task:

1. In the **Explorer** menu, choose **Search**.

   The navigation area displays items for constructing filter conditions.

2. Select the relevant items from the pull-down menus.
3. Click the **Search** button.

   The application area displays the Search subwindow, which contains the search results.

Stopping executed tasks

A migration task that is being executed or is pending execution can be stopped. For a task that is being executed, the processing in the storage subsystem to stop that task normally does not occur at the same time that the request to stop that task is issued. However, by setting the appropriate option, you can immediately stop migration tasks from being executed in the storage subsystem. Only migration tasks can be stopped in this way.

To stop a task that is being executed or is pending execution:

1. In the **Explorer** menu, choose **Tasks**.

   The Tasks subwindow appears in the application area.

2. Select the check box for the *task-ID* of each task to be stopped.
3. Click the **Stop Tasks** button.

   If at least one of the selected tasks is neither executing nor pending execution, the Stop Tasks dialog box appears, informing you which tasks can be stopped. In this case, proceed to step 4.

   If all the selected tasks can be stopped, a dialog box indicating that the tasks are being stopped appears. In this case, proceed to step 5.
4. Click the **OK** button. To immediately stop migration tasks being executed, select the **Stop tasks immediately only migration task** check box, and then click **OK**.

   A dialog box appears, indicating that tasks are being stopped. When the process is finished, a dialog box appears.
5. Click the **Close** button.

   The Stop Tasks dialog box closes and the Tasks subwindow re-appears showing the updated Tasks list.

Canceling tasks

To cancel a task that is on standby:

1. In the **Explorer** menu, choose **Tasks**.

   The Tasks subwindow appears in the application area.
2. Select the check box for the *task-ID* of each task that is to be canceled.
3. Click the **Cancel Tasks** button.

   If at least one of the selected tasks cannot be canceled, the Cancel Tasks dialog box appears, informing you which tasks can be canceled. In this case, proceed to step 4.

   If all the selected tasks can be canceled, a dialog box indicating that the tasks are being canceled appears. In this case, proceed to step 5.
4. Click the **OK** button.

   A dialog box indicating that the tasks are being canceled appears. When the cancellation process is finished, a dialog box appears.
5. Click the **Close** button.

   The Cancel Tasks dialog box closes and the Tasks subwindow re-appears showing the updated Tasks list.

Deleting tasks

You can delete completed tasks (tasks whose status is `Success`, `Failure`, `Cancel`, or `Stop`).

To delete a completed task:

1. In the **Explorer** menu, choose **Tasks**.

   The Tasks subwindow appears in the application area.
2. Select the check box for the *task-ID* of each task that is to be deleted.
3. Click the **Delete Tasks** button.

The Delete Tasks dialog box appears, confirming the tasks that are to be deleted.

4. In the Delete Tasks dialog box, select the **Yes. I have read the above, and still want to continue** check box.
5. Click the **OK** button.

   A dialog box appears, indicating that the tasks are being deleted. When the deletion process is finished, a dialog box appears.
6. Click the **Close** button.

   The Delete Tasks dialog box closes and the Tasks subwindow re-appears showing the updated Tasks list.

Editing task properties

To edit the properties of a task that has not ended:

1. In the **Explorer** menu, choose **Tasks**.

   The Tasks subwindow appears in the application area.
2. Click the  icon for the *task-ID* to be edited.

   The Edit Task dialog box appears.
3. Edit the properties of the task.

   You can edit the Description and Notification Email attributes.
4. Click the **OK** button.

4

Troubleshooting

This chapter identifies errors that might occur when using the Tiered Storage Manager Web client, describes their possible causes, and provides recommended corrective actions. If an error occurs, refer to the message manual and log, and then perform the appropriate recovery procedures.

- Responding to error conditions
- Troubleshooting on the Web client
- Viewing log data

Responding to error conditions

If an error occurs in Tiered Storage Manager, follow the flow of procedures shown in the figure below and take the appropriate action.

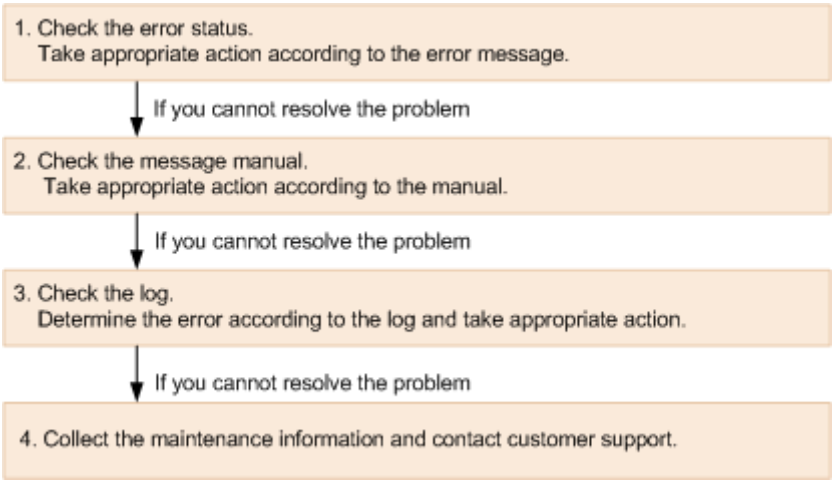

Figure 4-1: Steps to take when an error occurs

To fix the problem:

1. Check the error status.

   Check the status at the time the error occurred. If a message has been displayed, take the appropriate action. If you cannot resolve the problem, proceed to step 2.

2. Check the message manual.

   Take action based on the measures described in the message manual. For details on reasons and measures for messages output by Hitachi Storage Command Suite Common Component or by the CLI client, Web client, or server of Tiered Storage Manager, see *Hitachi Tiered Storage Manager Messages*. If you cannot resolve the problem, proceed to step 3.

3. Locate the error log, examine the log data for the Tiered Storage Manager server, and then take corrective action.

   Using the message logs, identify where the error occurred, and then remove the cause of the error. Use the message ID to locate the error in the log. For details on the message ID, see the Hitachi Tiered Storage Manager Messages. If you cannot resolve the problem, proceed to step 4.

4. Contact customer support.

   Use the tool for collecting maintenance information, and then contact Customer Support. For details about the tool for collecting maintenance information, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

Troubleshooting on the Web client

This section shows examples of troubleshooting for each function provided by Tiered Storage Manager.

Table 4-1: Login problems

Problem (displayed message ID)	Cause	Action
Unable to log in. (KATS20505-E)	The user is not registered.	Log in as a user with User Management permission, and then register the desired user.
	You do not have permission to log in to Tiered Storage Manager.	Log in as a user with User Management permission, and then assign the desired Tiered Storage Manager permissions to the user.
	The user ID or password you entered is incorrect.	Enter a correct user ID and password.
	The user account is locked.	Manually unlock the user account.

Table 4-2: Window display problems

Problem	Cause	Action
In an English environment, messages are displayed in a foreign language.	A foreign locale is set on the management server.	Specify an English locale on the management server.#1
After performing an overwrite installation of Hitachi Tiered Storage Manager, windows are not displayed correctly.	The style sheet used by Tiered Storage Manager has been updated, or the old style sheet was read from the cache.	Take the following action. In Internet Explorer: Select **View**, and then **Refresh**. Alternatively, select **Tools** and then **Internet Options**. Then, click **Delete Files** from **Temporary Internet Files**. In Mozilla or Firefox: Select **View**, and then **Reload**.
Dialog boxes are not displayed.	Popup windows are being blocked by Internet Explorer for Windows XP Service Pack 2 or later.	Register information about the target server as follows: 1. In Internet Explorer, choose **Tools**, **Pop-up Blocker**, and then **Pop-up Blocker Settings**. 2. In **Address of Web site to allow**, enter the address (URL) of the target server, and then click **Add**. When registration is successful, the server is added to **Allowed sites:**

Problem	Cause	Action
In the Windows environment, items in Explorer, the navigation area, or the application area do not display correctly.	Some items in the Internet Explorer Security Settings dialog box have been set to **Disable**.	Take the following action. 1. In Internet Explorer, choose **Tools**, and then **Internet Options**. In the **Security** tab, click the **Custom Level** button. 2. **Make** sure the following security settings are not set to **Disable**. **ActiveX controls and plug-ins** • **Run ActiveX controls and plug-ins** • **Initialize and script ActiveX controls not marked as safe** **Scripting** • **Active scripting** **Miscellaneous** • **Launching programs and files in an IFRAME** • **Submit not encrypted form data**
In a Solaris environment, characters in the browser title bar are not displayed correctly.	The browser language setting does not match the language set in the X-terminal.	Reset the browser language setting so that it matches the language set in the X-terminal.
After a database is restored, the information displayed in the navigation area and application area do not match.	The database of the Hitachi Storage Command Suite Common Component has not been synchronized with the database of Tiered Storage Manager.	Set `server.base.initialsynchro` in the `server.properties` file[#2] of Tiered Storage Manager to `true`, and then restart the server.

#1
An error message might contain text that is displayed as it was received from the Device Manager server. The language for these parts of a message is determined by the locale of the computers on which the Device Manager server and Tiered Storage Manager servers run. Therefore, we recommend that you use the same language type for the Web client and servers.

#2
For details on the `server.properties` file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

Table 4-3: Problems with Tiered Storage Manager functions

Problem (displayed message ID)	Cause	Actions
In Step 3/6 of the Migration Wizard, no compatible target storage tiers are listed.	There are no storage tiers that include enough migration targets for the volumes in the selected migration group.	Tiered Storage Manager determines that a storage tier is compatible as a migration target if it includes volumes that satisfy all of the following conditions: • For each volume in the migration group, the storage tier includes at least one volume that is of equal or greater capacity. • The emulation type of each target volume is the same as that of its source volume. (*Exception*: If a migration is performed from a CVS volume to a normal volume, the emulation type of both volumes must be OPEN-V.[#]) • The target volume is unused. (*Exception*: If the source volume is in the target storage tier, Tiered Storage Manager automatically selects the source volume as the target volume.) If this problem occurs, prepare a compatible storage tier based on the above conditions.

Problem (displayed message ID)	Cause	Actions
The progress of a task to migrate or erase data does not change.	Possible causes are as follows: • An error might have occurred in the subsystem and an attempt to acquire the progress status has failed. • Processing to update the display has not started because multiple migration tasks are being performed and processing to start up or finish these tasks might be taking a long time. • If a high load is placed on the subsystem, migration might not progress or might fail before completion. In particular, if a volume that is being migrated is accessed, progress of the migration task might deteriorate.	• If you believe that an error occurred in the subsystem, check the error messages, and remove the cause of any errors. • If at least one task among multiple migration tasks is progressing, Tiered Storage Manager is running properly. • If migration fails, temporarily decrease the load on the subsystem, and then retry the migration.
In Step 5/6 of the Migration Wizard#, clicking the **Confirm** button resulted in an error. (KATS62010-E or KATS62011-E)	Possible causes are as follows: • There is a problem with a volume that was selected as a migration target. • A selected volume cannot be reserved for migration because Storage Navigator is used in Modify mode. The details of this error are displayed in the Error window.	Remove the cause of the error, and then re-create the migration task. Depending on the timing of the error, some volumes might have been reserved. In this case, use Tiered Storage Manager to refresh the storage domain containing the reserved volumes.

Problem (displayed message ID)	Cause	Actions
The status of a migration task is **Eraser Failure**. (KATS62010-E or KATS62011-E)	An error occurred while erasing data. The details of the error are displayed in the Tasks subwindow.	Migration has finished and the LDEVs of the migration source volume and the migration target volume have been switched. Check the status of the source volume whose LDEV has been switched to the LDEV number of the target volume. If data erasure failed, the original data remains in the source volume. Therefore, perform one of the following: • If the volume to which the LDEV number of the target volume has been switched is blocked Format the relevant volume by using Storage Navigator. • If the volume to which the LDEV number of the target volume has been switched is not blocked There might be data remaining on the relevant volume. If the remaining data needs to be erased, format the volume or erase the data by using the procedure below. In the explanation below, LDEV 10:01 has been migrated to 20:01 (after migration, the LDEV of the source volume was switched to 20:01). 1. Assign a path from a host to LDEV 20:01. 2. From the above host, overwrite the data on LDEV 20:01 with 0 data. 3. Remove the path to LDEV 20:01.

Problem (displayed message ID)	Cause	Actions
The status of a migration task is **Failure**. (KATS62010-E or KATS62011-E)	An error occurred during migration. The details of the error are displayed in the Tasks subwindow.	Depending on the cause of the error, migration might be treated as a migration failure in Tiered Storage Manager or Device Manager, even though migration has finished in the storage subsystem. In this case, perform the following: 1. Refresh the relevant storage domain by using Tiered Storage Manager to update the management information of Tiered Storage Manager and Device Manager. If a volume remains reserved, it will be released during the refresh operation. 2. View the volume information of migration tasks that have the **Failure** status to determine whether the migration task has been executed. At this time, check the following elements for all LDEVs in the migration group: - Array group name - Subsystem name If there is a logical device for which migration has not been executed, remove the cause of the error, and then create and execute the migration task again.
Shredding is disabled after migration.	Storage Navigator was used in the external storage subsystem to lock a volume.	Unlock the volumes in the external storage subsystem, and then perform shredding.
Executing a migration task does not achieve migration.	Storage Navigator was used in the external storage subsystem to lock a volume.	Unlock the volumes in the external storage subsystem, and then perform migration.
When a scheduled event occurs, no email is sent.	Possible causes: • An SMTP server has not been set up. • SMTP authentication failed. • The specified email address is invalid. • Mail-related properties in server.properties are not correct.	• Make sure an SMTP server is connected. • Make sure SMTP authentication user information is correct. • Make sure the email address is correct. • Make sure the mail-related properties in server.properties are correct.

Problem (displayed message ID)	Cause	Actions
# If the microcode version of Universal Storage Platform V is 60-03-20-00/00-00 or later and both the migration source volumes and target volumes are DP volumes, or the migration source volume is a DP volume and the migration target volume is a DP pool, migration is possible between volumes whose emulation types are OPEN-V and OPEN-0V.		

Viewing log data

The Web client outputs the following types of log data on the Tiered Storage Manager server computer:

- Message logs
- Trace logs

Message logs

The message log contains a history of executed processing and processing results. The following table lists the name and output destination of the message log files:

Table 4-4: Name and output destination of message log files

Type of log information	Log file name	Output destination folder or directory
GUI message log	`HTSMGuiMessage`*n*`.log`#	In Windows *Tiered-Storage-Manager-installation-folder*`\logs` In Solaris or Linux `/var/opt/HiCommand/TieredStorageManager/logs`
# *n* represents a number from `1` to the *number-of-files-specified-in-the-environment-settings*. To specify the number of files, change the value of the `logger.guiMessageFileCount` in the environment settings file `logger.properties`. For details about this file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.		

Trace logs

The trace log contains information collected when errors occur that can be used to help determine the cause of the error. The following table lists the name and output destination of the trace log files:

Table 4-5: Name and output destination of trace log files

Type of log information	Log file name	Output destination folder or directory
GUI trace log	`HTSMGuiTrace`*n*`.log`#	In Windows *Tiered-Storage-Manager-installation-folder*`\logs` In Solaris or Linux `/var/opt/HiCommand/TieredStorageManager/logs`

#

n represents a number from `1` to the *number-of-files-specified-in-the-environment-settings*. To specify the number of files, change the value of `logger.guiTraceFileCount` in the environment settings file `logger.properties`. For details about this file, see the *Hitachi Tiered Storage Manager Server Configuration and Operation Guide*.

Reference Information

Tiered Storage Manager contains numerous functions and operates within a complex and varied environment. This appendix contains detailed listings of various properties, values and parameters that can be useful in working with the Tiered Storage Manager GUI.

- Supported storage subsystems
- Volumes as sources or targets
- Volume filter conditions
- Pool filter conditions
- Suggested maximum operational usage

Supported storage subsystems

This section describes the storage subsystems that can be used with Tiered Storage Manager and their supported firmware versions. The functions available in Tiered Storage Manager differ depending on the supported firmware version of the storage subsystem.

Storage subsystems that can be registered as domain controllers

Multiple domain controllers can be registered in Tiered Storage Manager. Table A-1:Subsystems that can be used as domain controllers lists the storage subsystems that can be used as domain controllers and the supported firmware versions.

Table A-1: Subsystems that can be used as domain controllers

Product name	Supported firmware version[#]
Hitachi Universal Storage Platform V	• DKC microcode 60-01-24 or later • DKC microcode for using Hitachi ShadowImage for Mainframe or Hitachi TrueCopy for Mainframe 60-01-62 or later • DKC microcode for migrating DP volumes: 60-02-20-00100 or later For details about migrating DP volumes, see Migrating data to DP pools. • DKC microcode for migration between OPEN-V and OPEN-0V: 60-03-20-00/00-00 or later For details about migration between OPEN-V and OPEN-0V, see Volume availability for migration. • DKC microcode when a volume linked to a TrueCopy pair is used as the migration source and an external volume is specified as the migration target: 60-03-27-00/00 or later. For details, see Volumes whose usability as migration sources or targets can be determined. • DKC microcode for using Universal Replicator or Hitachi Universal Replicator for Mainframe 60-03-28 or later • DKC microcode when a volume linked to a Universal Replicator pair is used as the migration source and an external volume is specified as the migration target: 60-04-00-00 or later For details, see Volumes whose usability as migration sources or targets cannot be determined. • DKC microcode for creating and executing an External Mapping Task: 60-04-00/00 or later For details about external mapping tasks, see External mapping tasks DKC microcode for using volumes that make up a LUSE 60-04-10 or later • DKC microcode when a volume linked to a TrueCopy pair with the status PAIR or Duplex is used: 60-05-00-00/00 or later For details, see Volumes whose usability as migration sources or targets can be determined.

Product name	Supported firmware version#
Hitachi Universal Storage Platform VM	• DKC microcode 60-01-24 or later • DKC microcode for using Hitachi ShadowImage for Mainframe or Hitachi TrueCopy for Mainframe 60-01-62 or later • DKC microcode when a volume linked to a TrueCopy pair is used as the migration source and an external volume is specified as the migration target: 60-03-27-00/00 or later For details, see Volumes whose usability as migration sources or targets can be determined. • DKC microcode when a volume linked to a Universal Replicator pair is used as the migration source and an external volume is specified as the migration target: 60-04-00-00/00 or later For details, see Volumes whose usability as migration sources or targets can be determined. • DKC microcode for creating and executing an External Mapping Task: 60-04-00 or later For details about the External Mapping Task, see External mapping tasks. DKC microcode for using volumes that make up a LUSE 60-04-10 or later • DKC microcode when a volume with the status of PAIR or Duplex is part of a TrueCopy pair is used as the migration source: 60-05-00-00/00 or later For details, see Volumes whose usability as migration sources or targets cannot be determined.
Hitachi USP	• DKC microcode 50-03-93 or later 50-04-48 or later 50-05-22 or later 50-06-00 or later 50-07-00 or later • DKC microcode for using Universal Replicator 50-09-85 or later

External storage subsystems

The following external storage subsystems can be managed by Tiered Storage Manager:

- Lightning 9900V
- Thunder 9500V
- Lightning 9900
- Hitachi AMS series
- Hitachi AMS2000 series
- Hitachi SMS series
- Hitachi WMS series

- Hitachi USP
- Universal Storage Platform V
- Universal Storage Platform VM

In addition to the storage subsystems listed above, Tiered Storage Manager can also manage storage subsystems supported by the Universal Storage Platform V/VM external storage connection function or the Hitachi USP external storage connection function. Note, however, that when a storage subsystem not supported by Device Manager (which is a required program) is used as an external storage subsystem, volume attributes might be displayed as **Unknown**.

Table A-2:SMI-S Enabled subsystems that can be used in Tiered Storage Manager shows the SMI-S Enabled subsystems that can be used in Tiered Storage Manager.

Table A-2: SMI-S Enabled subsystems that can be used in Tiered Storage Manager

Vendor	Model	Required SMI-S provider
HP	• EVA 3000#1 • EVA 4100#1 • EVA 4400#1 • EVA 5000 • EVA 6100#1 • EVA 6400#1#2 • EVA 8100#1 • EVA 8400#1#2	• HP Command View EVA 8.0.2 or later (SMI-S: Version 1.2 or later)
EMC	• CLARiiON CX200 • CLARiiON CX300 • CLARiiON CX400 • CLARiiON CX500 • CLARiiON CX600 • CLARiiON CX700 • CLARiiON CX3 series	• EMC SMI-S Provider 3.3 or later (prerequisite program: EMC Solutions Enabler 6.5 or later) (SMI-S: Version 1.2 or later)
#1: This model can be used when the DKC microcode is 5.031 or later. #2: The required microcode version of USP V/VM is 60-03-10-xx/xx or later.		

In EMC CLARiiON CX series, volumes created in RAID level 0 cannot be mapped externally.

We recommend that you make the interval at which storage subsystem information managed by the SMI-S provider is reflected in the SMI-S provider as short as possible.

Storage subsystems in which volumes can be created

Tiered Storage Manager can create volumes in the following storage subsystems:

- Lightning 9900V
- Thunder 9500V
- Lightning 9900
- Hitachi AMS series
- Hitachi AMS2000 series
- Hitachi SMS series
- Hitachi WMS series
- Hitachi USP
- Universal Storage Platform V
- Universal Storage Platform VM

In addition to the storage subsystems listed above, volumes can be created in the SMI-S Enabled subsystems listed in Table A-2:SMI-S Enabled subsystems that can be used in Tiered Storage Manager.

Volumes as sources or targets

In Tiered Storage Manager, some volumes cannot be used as migration sources or migration targets. This section describes the volumes that can and cannot be used in Tiered Storage Manager.

For Tiered Storage Manager, volumes fall into one of two categories:

- Volumes whose usability as migration sources or targets can be determined
- Volumes whose usability as migration sources or targets cannot be determined

Volumes whose usability as migration sources or targets can be determined

If Tiered Storage Manager determines that a volume cannot be used as a migration source or migration target, **No** is displayed in **Can Migrate**, and the reason that the volume cannot be used is displayed in **Migration Restrictions**. Table A-3:Volumes that cannot be used as migration sources or migration targets describes the volumes that cannot be used as migration sources or migration targets.

Table A-3: Volumes that cannot be used as migration sources or migration targets

A volume cannot be used as a migration source or target, if...	Character string displayed in Migration Restrictions
Either of the following is set for the specified volume: • Cache Residency Manager • Cache Residency Manager for z/OS®	**Cache Residency Manager**
The specified volume is used as a command device.	**Command Device**
One of the following applies to the volume: • The volume is the P-VOL of a Copy-On-Write Snapshot pair, its status is `Pair`, and the number of generations is 64. • The volume is the P-VOL of a Copy-On-Write Snapshot pair and its status is not `Pair`. • The volume is the V-VOL of a Copy-On-Write Snapshot pair. • The volume is the POOL of a Copy-On-Write Snapshot pair.	**Copy-On-Write Snapshot Status**
The specified volume is a Dynamic Provisioning pool volume.	**Dynamic Provisioning [Pool]**
The specified volume is a Dynamic Provisioning virtual volume and is not associated with a pool.	**Dynamic Provisioning [Unassigned]**
The volume is using the external subsystem connection function, which is locked by Data Retention Utility (is not Read/Write).	**Externally Locked**
The specified volume has already been reserved as a migration target.	**Migration Reserved**[#]
Information for the specified volume cannot be acquired.	**Not Acquired Volume Information**
Either of the following applies to the specified volume: • The volume is a ShadowImage S-VOL, and is paired with an SP-VOL. • The volume is a Hitachi ShadowImage for Mainframe S-VOL, and is paired with an SP-VOL.	**ShadowImage Configuration [Leaf]**
Either of the following applies to the specified volume: • The volume is a ShadowImage SP-VOL, and is paired with two S-VOLs. • The volume is a Hitachi ShadowImage for Mainframe SP-VOL, and is paired with two S-VOLs.	**ShadowImage Configuration [Node]**

A volume cannot be used as a migration source or target, if...	Character string displayed in Migration Restrictions
Either of the following applies to the specified volume: • The volume is a ShadowImage P-VOL, and is paired with three S-VOLs. • The volume is a Hitachi ShadowImage for Mainframe P-VOL, and is paired with three S-VOLs.	**ShadowImage Configuration [Root]**
The specified volume is a system disk.	**System Disk**
Either of the following applies to the specified volume: • The volume is a part of a TrueCopy pair, and the status is neither `PSUS` nor `PSUE`. • The volume is a part of a Hitachi TrueCopy for Mainframe pair, and the status is not suspended.	**TrueCopy Asynchronous Status**
Either of the following applies to the specified volume: • The volume is a part of a TrueCopy pair, and the status is neither `PSUS`, `PSUE`, or `PAIR`. • The volume is a part of a Hitachi TrueCopy for Mainframe pair, and the status is not suspended.	**TrueCopy Synchronous Status**
The volume is in a Hitachi USP, and either of the following applies: • The volume is a part of a TrueCopy pair, and the status is `PAIR`. • The volume is a part of a TrueCopy for Mainframe pair, and the status is `Duplex`.	**TrueCopy Asynchronous Status** **TrueCopy Synchronous Status**
One of the following applies to the specified volume: • The volume is a part of a Universal Replicator pair, and the status is `COPY` or `PAIR`. • The volume is a part of a Hitachi Universal Replicator for Mainframe pair, and the status is `Pending duplex` or `Duplex`. • The volume is a Universal Replicator journal volume. • The volume is a Hitachi Universal Replicator for Mainframe journal volume.	**Universal Replicator Status**
# Includes the volumes migrated by Performance Monitor.	

When Tiered Storage Manager determines that a volume can be used as a migration source but cannot be used as a migration target, **Yes** is displayed in **Can Migrate**, and **(target)** is appended to the character string displayed in **Migration Restrictions**. Table A-4:Volumes that can be used as migration sources but not as migration targets describes the volumes that cannot be used as migration targets.

Table A-4: Volumes that can be used as migration sources but not as migration targets

Reason a volume cannot be used as a migration target	Character string displayed in Migration Restrictions
The volume is the P-VOL of a Copy-On-Write Snapshot pair, its status is `Pair`, and the number of generations is 63 or less.	**Copy-On-Write Snapshot (target)**
The specified volume is locked by the Data Retention Utility (not Read/Write).	**Locked (target)**
The specified volume is part of a LUSE	**LUSE (target)**
The specified volume is included in another migration group.	**Migration Group (target)**
The specified volume is a NAS system volume.	**NAS (target)**
A path is set for the specified volume.	**Path (target)**
Either of the following applies to the specified volume: • The volume is a part of a ShadowImage pair. • The volume is a part of a Hitachi ShadowImage for Mainframe pair.	**ShadowImage (target)**
Either of the following applies to the specified volume: • The volume is a part of a TrueCopy pair, and the status is `PSUS` or `PSUE`. • The volume is a part of a Hitachi TrueCopy for Mainframe pair, and the status is `suspended`.	**TrueCopy Asynchronous (target)**
One of the following applies to the volume: • The volume is a part of a TrueCopy pair, and the status is `PSUS` or `PSUE`. • The volume is a part of a TrueCopy pair on Universal Storage Platform V/VM, and the status is `PAIR`. • The volume is a part of a Hitachi TrueCopy for Mainframe pair, and the status is `suspended`. • The volume is a part of a Hitachi TrueCopy for Mainframe pair on Universal Storage Platform V/VM, and the status is `Duplex`.	**TrueCopy Synchronous (target)**
Either of the following applies to the specified volume: • The volume is a part of a Universal Replicator pair for z/OS, and the status is neither `COPY` nor `PAIR`. • The volume is a part of a Hitachi Universal Replicator for Mainframe pair, and the status is neither `Pending duplex` nor `Duplex`.	**Universal Replicator (target)**
VOLSER is set for the specified volume.	**VOLSER (target)**

When Tiered Storage Manager determines whether a volume can be used as a migration source or target, if a volume is registered in a migration or shredding task, **No** is displayed for **Can Migrate**, but no reason is displayed in **Migration Restrictions**. In addition, the user cannot use such volumes as migration source volumes or migration target volumes.

Volumes whose usability as migration sources or targets cannot be determined

Tiered Storage Manager is sometimes unable to determine whether some volumes can be used as migration sources or targets. Accordingly, a volume might not be suitable as a migration source or target even if **Yes** is displayed for **Can Migrate** and **Free** is displayed for **Volume Status**. For such volumes, the reason that migration is not possible is not displayed in **Migration Restrictions**. If such volumes are selected as migration sources, an error occurs when the migration task is executed. If such volumes are selected as migration targets, an error occurs when the migration task is created. The volumes that satisfy the conditions listed below are in this category. We recommend that you register any such volumes in a migration group whose attribute disables migration to prevent them from being selected as migration sources or migration targets.

- Volumes that are part of a Cross-System Copy pair
- Volumes that are configured for XRC (IBM's remote copy facility, Extended Remote Copy)

  If the status is suspended, the volume can be used as the migration source volume.
- Volumes that are configured for concurrent copying
- Volumes that cannot be accessed, because they are in blocked status or being formatted
- Volumes on which data migration is being performed by Performance Manager
- Volumes that have been set by an optional program for mainframe systems (such as Hitachi TrueCopy Synchronous for Mainframe)
- Migration target DP pool volumes that consist of volumes on different CLPRs and SLPRs.
- The migration source volume is configured as a TrueCopy pair or Hitachi TrueCopy for Mainframe pair, and is included in an external storage subsystem.
- The migration source volume is configured as a Universal Replicator pair or Hitachi Universal Replicator for Mainframe pair, and is included in an external storage subsystem.
- The migration source is configured as a TrueCopy Synchronous or TrueCopy Asynchronous pair, and the migration target is a DP volume.
- The migration source volume is configured as a Universal Replicator pair, and the migration target volume is a DP volume.
- The migration source volume is the S-VOL of a Universal Replicator pair for Delta Resync.

- The migration source volume is the S-VOL of a Hitachi Universal Replicator for Mainframe pair for Delta Resync.
- The volume is reserved for a ShadowImage or Hitachi ShadowImage for Mainframe operation.
- The volume is being quick-formatted.
- The volume is being shredded.
- This volume is set in Volume Security as unable to be used as an S-VOL.
- The volume is configured as a Compatible XRC pair.

  If the status is `suspended`, the volume can be used as a migration source.

When any of the following conditions apply, an error occurs during the execution of a migration task:

- For a Universal Storage Platform V/VM subsystem whose microcode version is earlier than 60-03-27-00, a volume linked to a TrueCopy Synchronous pair or TrueCopy Asynchronous pair is specified as the migration source volume, and an external volume is specified as the migration target volume
- For a Universal Storage Platform V/VM subsystem whose microcode version is earlier than 60-03-27-00, a volume linked to a Hitachi TrueCopy Synchronous for Mainframe pair or a Hitachi TrueCopy Asynchronous for Mainframe pair is specified as the migration source volume, and an external volume is specified as the migration target volume
- For a Universal Storage Platform V/VM subsystem whose microcode version is earlier than 60-04-00-00, a volume linked to a Universal Replicator pair is specified as the migration source volume, and an external volume is specified as the migration target volume.
- For a Universal Storage Platform V/VM subsystem whose microcode version is earlier than 60-04-00-00, a volume linked to a Hitachi Universal Replicator for Mainframe pair is specified as the migration source volume, and an external volume is specified as the migration target volume.

Volume filter conditions

When searching for volumes, the filter conditions available can vary according to the various selections made.

Filter conditions available for existing tiers and domains

The following table lists filter conditions for use in identifying volumes:

- For migration to a storage tier that has already been created.
- During volume search, when a storage domain name is already specified for **Search In**.

Table A-5: Filter search conditions for existing tiers and domains

Filter Condition	Description
Device Number	Specifies the logical device number of the controller. Enter two or three hexadecimal strings using the following rules: • Each string should be less than 0x100, separated with a colon (:). • Values separated by a comma (,) and ranges separated by a hyphen (-) can be used, but not together. • A space cannot be specified before or after a colon. A space used before or after a comma or hyphen is disregarded. Only one range can be specified. • As many as 100 elements can be specified. • Alphanumeric characters can be specified. Example: 0:0A, 0:0B, or 0:0A - 0:0C
I/O Consumer	Specifies the name of the usage source accessible to the volume on the domain controller. In this version, a host name can be specified for the usage source. • A string of 256 bytes or fewer after UTF8 conversion can be specified. • Multiple values cannot be specified. • Both one-byte and double-byte characters can be specified. • If the "starts with" operator is specified, any space specified at the end of the string becomes part of the search pattern. • If the "contains" operator is specified, any space specified at the beginning of the string becomes part of the search pattern. 1 character = 1 to 3 bytes 1 ASCII character = 1 byte
Volume Status	Specifies the usage status of the volume. • Specify "Used" or "Free" from the drop-down list.
Subsystem	Specifies the name of the subsystem. If the storage is not registered with Device Manager, the product name and serial number are displayed as "ProductName#SerialNumber".
Subsystem Vendor	Specifies the name of the vendor. The vendor name is displayed, such as "HITACHI" or "HP".
Subsystem Display Model	Specifies the subsystem model name to display. "Hitachi_USP", or other is displayed.
Subsystem Model	Specifies the subsystem model name: • During volume search, nothing is displayed. • During tier creation or modification, the subsystem model name is only displayed when the system was upgraded from one with the V4.x SubsystemModel attribute. "HDS9970V", "HDS9530V", or other is displayed. The model name is displayed for the subsystem model, not the model name for display. If the model name obtained from Device Manager is 'Unknown', the product name is displayed as "ProductName".
Subsystem Serial Number	Specifies the subsystem serial number.

Filter Condition	Description
Ctrl. Array Group	Specifies the array group on the controller. • A string of 150 bytes or fewer after UTF8 conversion can be specified. • Multiple values cannot be specified. One-byte and double-byte characters can be specified. • If the "starts with" operator is specified, any space specified at the end of the string becomes part of the search pattern. • If the "contains" operator is specified, any space specified at the beginning of the string becomes part of the search pattern.
Array Group	Specifies the array group. • See the online help.
Array Group Busy Rate#	Specifies the array group usage. • Enter an integer from "0 to 100". One-byte numerals can be specified.
Array Group Max Busy Rate#	Specifies the maximum value for array group usage. • Enter an integer from "0 to 100". One-byte numerals can be specified.
Array Group Encryption	Specifies the array group encryption attribute. • See the online help.
RAID Level	Specifies the RAID level. This can also be specified with only the RAID Level. • Specify one of the following from the drop-down list: "RAID0", "RAID0+1", "RAID1", "RAID1+0", "RAID5", "RAID6" Specifies the drive configuration. This cannot be specified with only the drive configuration. • Specify one item from the drop-down list: "5D+1P", etc.
Disk Type	Specifies the disk type. Specify either "FC", "AT", "BD", "HD-FC", "-", "SAS", or "SSD"
Capacity	Specifies the volume capacity. • Enter an integer from 0 to 2^{63}-1. One-byte numerals can be specified. Specifies the units for volume capacity. • Specify one of the following from the drop-down list: "Block", "KB", "MB", "GB", "TB"
Volume Lock Status	Specifies the volume lock status. • Specify one of the following from the drop-down list: "Locked", "Unlocked"
Emulation Type	Specifies the volume emulation type. • The values displayed for the LogicalUnit object emulation attribute in GetStorageArray (subtarget=LogicalUnit), for HDvM CLI.
SLPR	Specifies the SLPR number. The CLPR number specification is optional. • Specify an integer from 0-31. One-byte numerals can be specified.
CLPR	Specifies the CLPR number. The SLPR number specification is optional. • Specify an integer from 0-31. One-byte numerals can be specified.
SYSPLEXID/DEVN	Specifies SYSPLEXID and DEVN, separated by a slash (/).

Filter Condition	Description
VOLSER	Specifies VOLSER.
Logical Group	Specifies the logical group. • A string of 1,536 bytes or fewer after UTF-8 conversion can be specified in the text box. HTSM only handles the first 1,024 characters in a path-format logical group name, up to 1,536 bytes after UTF-8 conversion. • To search the logical groups cut off in the latter half, the cut-off logical name needs to be specified.
Port/HSD	Specifies the port and host storage domain, separated with a slash (/). • A string of 149 bytes or fewer after UTF-8 conversion can be specified in the text box.
Disk RPM	Specifies the disk revolutions per minute (rpm). • Integers from 0 to 2147483645 can be specified.
Disk Capacity	Specifies the disk capacity (in GB). • Integers from 0 to 2147483645 can be specified.
P-VOL's Migration Group	Specifies the migration group for the PVOL.
P-VOL's MU Number	Specifies the MU number for the PVOL. • Integers from 0 to 2 can be specified.
ShadowImage	Specifies the ShadowImage volume type. • Specify one of the following from the drop-down list: "Simplex", "P-VOL", "S-VOL", "SP-VOL"
TrueCopy Synchronous	Specifies the TrueCopy Synchronous volume type. • Specify one of the following from the drop-down list:"Simplex", "P-VOL", "S-VOL"
TrueCopy Asynchronous	Specifies the TrueCopy Asynchronous volume type. • Specify one of the following from the drop-down list:"Simplex", "P-VOL", "S-VOL"
Universal Replicator	Specifies the Universal Replicator volume type. • Specify one of the following from the drop-down list:"Simplex", "P-VOL", "S-VOL", "SP-VOL", "JNL-VOL"
Copy-On-Write Snapshot	Specifies the Copy-On-Write-Snapshot volume type. • Specify one of the following from the drop-down list:"Simplex", "P-VOL", "V-VOL", "POOL"
CVS	Specifies the whether the volume is a CVS volume. • Specify one of the following from the drop-down list:"Yes", "No"
Dynamic Provisioning	Specifies the volume type. Specify one of the following from the drop-down list: • For versions earlier than 05.90: "-", "V-VOL", "POOL" • For version 05.90 or later: "-" "DP-VOL", "DP-Pool" • For version 06.00 or later: "-" "DP-VOL", "DP-Pool-VOL"
Consumed Capacity	Specifies the volume capacity used. • Values converted into KB can be from 0 to 281474976710653. Specifies the units for the volume capacity used. • Specify one of the following from the drop-down list:"KB", "MB", "GB", "TB"

Filter Condition	Description
Consumed Capacity Percentage	Specifies the percentage of capacity used. • Specify an integer from 0 to 100.
Pool ID	Specifies the pool number. • There is no character limit. The maximum length for 1 element is 20 bytes after UTF-8 conversion. • As many as 100 elements can be specified, separated by commas (,).
Label	Specifies the volume label. • The maximum is 64 bytes. One label can be specified for each text box.
# When ArrayGroupBusyRate and ArrayGroupMaxBusyRate are specified, no hits will be found when volumes with unknown array group usage (where usage is blank) are searched.	

Filter conditions for specific subsystems

Table A-6: Filter conditions available when a specific subsystem is selected.

Filter Condition	Description (#1, #2)
Device Number	Specifies the logical device number of the volume. • See Table A-8:Input restrictions for device numbers and array groups
Label	Specifies the volume label. The maximum is 64 bytes. One label can be specified for each text box.
Capacity	Specifies the volume capacity. • Enter Integers from 0 to 2^{63}-1 can be specified. One-byte numerals can be specified. Specifies the units for volume capacity. • Specify one of the following from the drop-down list: "KB","MB","GB","TB"
Path	Specifies whether a path is set for the volume. • Specify either "Yes"or "No"
Array Group	Specifies the array group. • See Table A-8:Input restrictions for device numbers and array groups
RAID Level	Specifies the RAID level. This can also be specified with only the RAID Level. • Specify one of the following: "RAID0", "RAID0+1","RAID1", "RAID1+0","RAID5", "RAID6"
	Specifies the drive configuration. This cannot be specified with only the drive configuration. • Specify "5D+1P", etc. For third-party storage, "-" is displayed.

Filter Condition	Description (#1, #2)
Volume Status	Specifies the volume usage status. • For Hitachi storage, specify one of the following: "Unknown", "Normal", "Blocked", "Format", "Correction Access", "Copying", "ReadOnly", "Shredding", "Normal (Quick Format)", "Preparing Quick Format", "Unformat", "Regression", "Detached", "Undefined", "Format", "FormatWaiting" • For third-party storage, specify one of the following: "Unknown", "OK", "Degraded/Warning", "Minor failure", "Major failure", "Critical failure", "Non-recoverable error"

#1:
An error will occur if nothing is specified in the text box.

#2:
Only items within the storage domain are displayed for the value. If a filter condition created by using the CLI does not exist on the storage domain, it will be displayed along with existing filter conditions.

Input rules for filter conditions

Table A-7: Items displayed for encryption attributes

Subsystem	Volume	Microcode	Encryption functionality	Encryption attribute display
USP_V/USP_VM	LUSE#2, DP-VOL, V-VOL	-	-	"-"
	Any volume other than the above	60-04-xx or later	Enabled (the array group is encrypted)	"Enabled"
			Disabled (the array group is not encrypted)	"Disabled"
		Earlier than 60-04-xx	Unsupported	"Disabled"
Any subsystem other than the above	-	-	Unsupported	"-"
Unknown encryption functionality#1	-	-	Unknown	"Unknown"

#1:
If an external subsystem undiscovered by HDvM is referenced, this is "Unknown".

#2:
If information about individual LDEVs comprising a LUSE is displayed (in the LDEV tab of the Volume Details window), each LDEV encryption attribute is displayed.

Table A-8: Input restrictions for device numbers and array groups

Item	Restriction	
	Using operator is, is not	**Using operator start with, contains**
Maximum string length across 1 device number or 1 array group	A string of 1,024 bytes or fewer after UTF-8 conversion can be specified.	A string of 1,024 bytes or fewer after UTF-8 conversion can be specified.
Multiple device number or multiple array group specifications, separated by comma	As many as 100 items.	Only 1 item.
Space before and after a device number name or an array group name	This cannot be specified. Note that any one-byte space before and after a comma in multiple specifications is disregarded.	This cannot be specified. Consider the following rules for specifying items using "start with" and "contains" operators: • If the "starts with" operator is specified, any space specified at the end of the string becomes part of the search pattern. • If the "contains" operator is specified, any space specified at the beginning of the string becomes part of the search pattern.
Characters that can be entered	There is no limit to the number of characters that can be entered. However, note the following input restrictions: • Enclose commas or single quotes specified in a device number name or an array group name in single quotes. • Escape any single quotes specified in a device number name. The following list shows examples that can be specified: • name • ‘name,2’ • ‘name’’3’ • ‘name’’,4’	

Pool filter conditions

Table A-9:Available pool search conditions describes the pool search conditions that are available in the Search window.

Table A-9: Available pool search conditions

Pool search condition	Description
Free Capacity	Specifies the free capacity of the pool. • Specify one of the following values "KB", "MB", "GB" (default), or"TB" and use '>', '>=', '<', '<=', '=', '<>' operators. Specifies the unit for the volume capacity.
Pool ID	Specifies the pool ID.
Over Provisioning Percent	Specify the Over Provisioning percent. • Enter an integer from 0 to 2147483645 in the text box, and use '>', '>=', '<', '<=', '=', '<>' operators.
Over Provisioning Warning	Specify the Over Provisioning Warning threshold. • Enter an integer from 0 to 2147483645 in the text box, and use '>', '>=', '<', '<=', '=', '<>' operators.
Over Provisioning Limit	Specify the Over Provisioning Limit threshold. • Enter an integer from 0 to 2147483645 in the text box, and use '>', '>=', '<', '<=', '=', '<>' operators.

Suggested maximum operational usage

For planning purposes, Table A-10:Suggested maximum values for Tiered Storage Manager operations contains suggested maximum values for Tiered Storage Manager operations. The suggested maximum values assume typical operations. These values should be used as guidelines and (except where noted) are not absolute maximum values.

Table A-10: Suggested maximum values for Tiered Storage Manager operations

Category	Property	Suggested maximum value for normal operations
Storage domains	Number of storage domains	5
	Number of external storage subsystems that can be connected to a single storage domain	5
Storage tiers	Number of storage tiers in a storage domain	100
Migration groups	Number of migration groups in a storage domain	5,000
Volumes	Number of volumes in a migration group	300
	Number of volumes in a storage tier	65,280
	Number of free volumes in a storage tier	1,300
	Number of volumes to which a path can be assigned at one time	100#

Category	Property	Suggested maximum value for normal operations
Tasks	Number of unfinished tasks that can be registered in Tiered Storage Manager	100
	Number of tasks that can be registered in Tiered Storage Manager	5,000
	Total number of volumes that can be included in all tasks registered in Tiered Storage Manager (the source and target volumes are counted separately)	30,000
Server execution multiplexing	Number of users who can access the Tiered Storage Manager server concurrently	5#
Volume filter conditions	Number of condition elements that can be combined by using AND or OR in one filter condition expression	30#
	Number of elements that can be concurrently specified for a single search condition when searching for LDEVs or array groups	100#

This is the maximum value. If this value is exceeded, an error occurs.

The following problems might occur if you use Tiered Storage Manager in an environment where the values used in normal operation are exceeded:

- The response time during Tiered Storage Manager operations becomes excessively long.
- A memory shortage occurs because the amount of memory used by the Web browser and Tiered Storage Manager server, GUI, and CLI processes increases.
- If response time is excessively long, an error (for example, KATS41055-E or KATS61005-E) occurs if another operation is attempted.

B

Details of the array group selection rules

This appendix describes how Tiered Storage Manager selects LDEVs based on the distribution method that is used. In the examples in this appendix, four migration target LDEVs are selected. These examples also assume that the array group avoidance rule is not specified for any of the cases.

- BalanceCapacity
- MaximumCoverage
- MinimumCoverage

BalanceCapacity

Tiered Storage Manager selects the migration target LDEVs so an equal distribution of LDEVs is maintained among the array groups in the target storage tier. Migration target LDEVs in array groups that have more LDEV candidates available as migration targets are selected before the LDEVs in array groups that have fewer LDEV candidates. If each array group has the same number of LDEV candidates available as migration targets, the LDEVs are selected in alphabetical order by controller array group name. In Figure B-1:Example with BalanceCapacity specified, array group B-B-B has six LDEVs that can be selected as migration targets, the largest number of the three array groups shown. Accordingly, three LDEVs are selected from array group B-B-B and one LDEV is selected from array group A-A-A. As a result, array groups A-A-A and B-B-B each have three remaining LDEVs that can be used as migration target candidates for other migration groups.

Figure B-1:Example with BalanceCapacity specified illustrates this example.

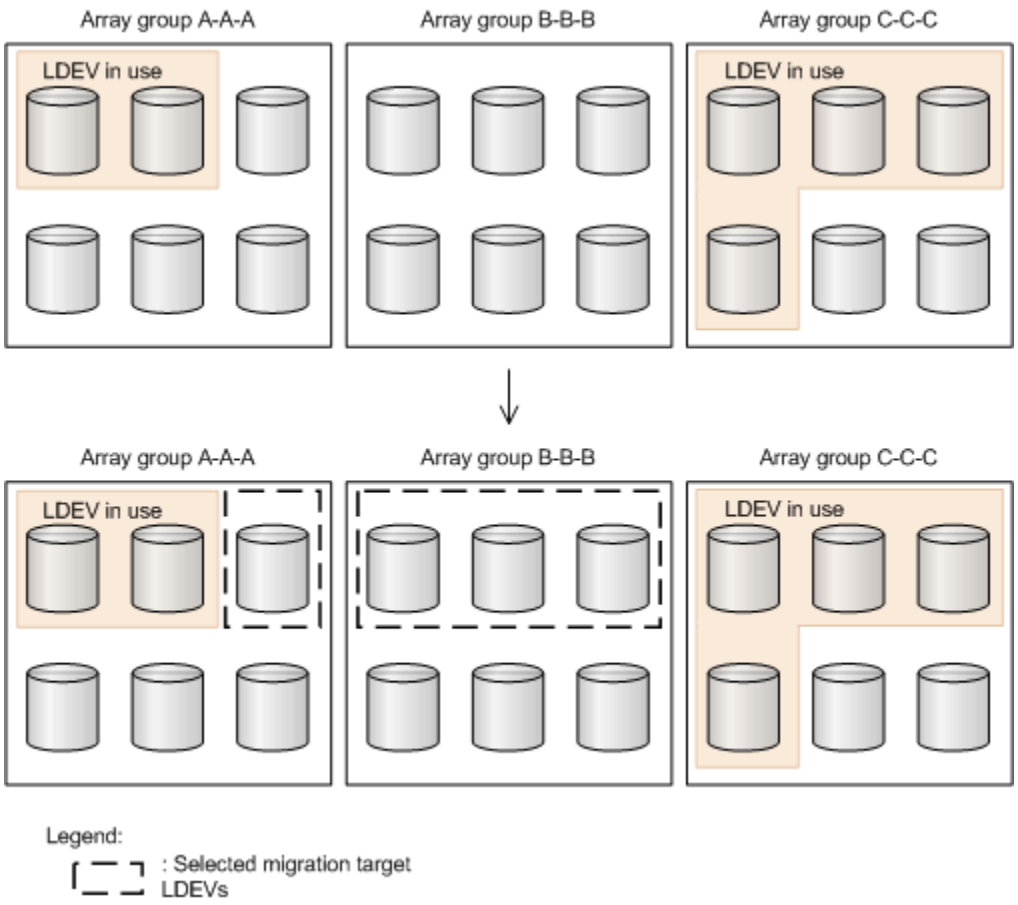

Figure B-1: Example with BalanceCapacity specified

MaximumCoverage

Tiered Storage Manager selects LDEVs so that the migration target LDEVs are distributed among as many array groups as possible. Tiered Storage Manager sorts the array groups based on the following conditions:

1. Array groups that do not contain any migration source LDEVs (array groups in which all LDEVs can be migrated)
2. Array groups that have the highest number of LDEVs available as migration targets
3. In alphabetical order by controller array group name

Tiered Storage Manager selects one LDEV from each array group in the order in which the array groups were sorted. When an LDEV has been selected from each of the array groups, Tiered Storage Manager returns to the first array group to select another LDEV. In Figure B-2:Example with MaximumCoverage specified, because array group B-B-B contains LDEVs that cannot to be migrated, LDEVs in this group are selected last. Of the two remaining array groups, array group A-A-A has more migration target LDEV candidates than array group C-C-C. Therefore, Tiered Storage Manager selects one migration target LDEV each from array group A-A-A, array group C-C-C, and then array group B-B-B, before returning to array group A A-A to select the last migration target LDEV. Figure B-2:Example with MaximumCoverage specified illustrates this example.

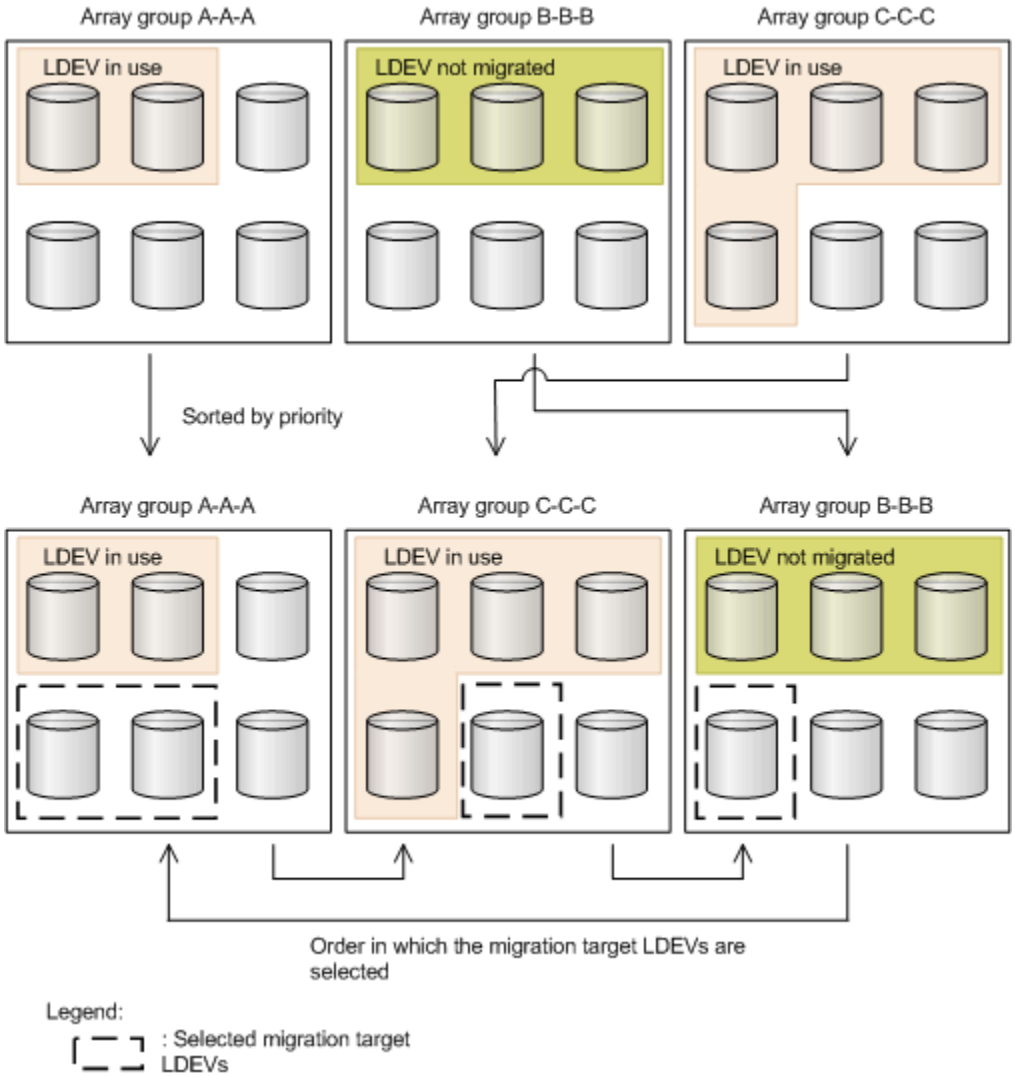

Figure B-2: Example with MaximumCoverage specified

MinimumCoverage

Tiered Storage Manager selects LDEVs so that they are distributed among as few array groups as possible. Tiered Storage Manager sorts the array groups based on the following conditions:

1. Array groups with LDEVs that cannot be migrated
2. Array groups that have the highest number of LDEVs available as migration targets
3. In alphabetical order by controller array group name

In Figure B-3:Example with MinimumCoverage specified, because array group B-B-B contains LDEVs that cannot to be migrated, LDEVs in this array group are selected first. Of the two remaining array groups, array group A-A-A has more migration target LDEV candidates than array group C-C-C. Therefore, the order of selecting migration target LDEVs is array group B-B-B, array group A-A-A, and then array group C-C-C. Because array group B-B-B has only three available LDEVs, one LDEV is also selected from array group A-A-A, which has the largest number of migration target LDEV candidates. Figure B-3:Example with MinimumCoverage specified illustrates this example.

Figure B-3: Example with MinimumCoverage specified

C

Relationship of terminology for different storage subsystems

The terminology used for storage subsystems can vary with hardware and software manufacturers. In general, the terminology used in this manual adheres to the definitions listed in the Storage Networking Industry Association (SNIA) dictionary. This chapter contains terminology specific to Tiered Storage Manager.

❒ Tiered Storage Manager terminology

Tiered Storage Manager terminology

Table C-1:Relationship of terminology for different storage subsystems lists the relationship between Hitachi storage subsystem terms and SNIA terms.

Table C-1: Relationship of terminology for different storage subsystems

No.	Hitachi Term	SNIA Term
1	Array Group	Storage Pool
2	DP Pool	
3	Snapshot Pool	
4	LDEV[#1]	Storage Extent
5	Free Space	
6	LU	Storage Volume
7	DevNum[#2]	DeviceID
8	SLPR	N/A
9	CLPR	
10	Emulation Type	
11	Command Device	SCSI Arbitrary Logical Unit
12	S-Vol	Source Volume
13	P-Vol	Target Volume
14	Subsystem[#3]	Storage Array
15	LUSE	Volume Composition
16	Data Retention Utility	Storage Element Protection
17	Host Storage Domain	SCSI Protocol Controller
18	WWN	Storage Hardware ID
19	Host Mode	Storage Client Setting Data
20	External Port	Initiator Port
21	ShadowImage	Local Replica (CopyType = Mirror)
22	TrueCopy (Sync)	Remote Replica (CopyType = Mirror), (Sync)
23	TrueCopy (Async)	Remote Replica (CopyType = Mirror), (Async)
24	CoW Snapshot	Delta Replica[#4]
25	PAIR (copy status)	Synchronized
26	PSUS (copy status)	Fractured
27	COPY(RS) (copy status)	Resync In Progress
28	N/A	Fracture In Progress[#5]

No.	Hitachi Term	SNIA Term
#1 *Storage Volume* corresponds to *LDEV* (which has the same meaning as *logical unit*) displayed in Tiered Storage Manager. #2 The format of the value also differs. The SCSI inquiry data appears in *DeviceID*. #3 `CIM_ComputerSystem` corresponds to the CIM class. #4 Might also be referred to as *Snapshot* when referring to an SMI-S Enabled subsystem. #5 The state of moving from a copy synchronous state to a split state.		

D

Notes about using earlier versions of the Web client

This appendix explains the differences between earlier versions of the Web client and the most recent version.

- Searching for tasks by specifying a storage subsystem name
- Volumes displayed as candidates for migration targets
- Message displayed when storage domain information must be refreshed
- The number of candidate volumes to be displayed
- Change in the default setting for deleting data
- Display of task owner names

Searching for tasks by specifying a storage subsystem name

In Tiered Storage Manager 6.1 or later, you can specify a storage subsystem name as a filter condition when searching for tasks. However, if you do this, only tasks created in Tiered Storage Manager 6.1 or later are displayed in the search results.

Volumes displayed as candidates for migration targets

In Tiered Storage Manager 6.0 or later, you can create a migration task even if the capacity of a target volume is greater than that of its source volume. This means that volumes that have a greater capacity than the migration source volume are also displayed as candidate volumes.

If you do not want volumes that have a greater capacity than a migration source volume to be displayed, specify `0` for the property `server.migrationPlan.candidateCapacityGroupDisplayMaxCount`.

Message displayed when storage domain information must be refreshed

In version 5.9 or later of the Tiered Storage Manager server, if the storage domain information has not been updated, a warning message is displayed, and you cannot create a task without refreshing the information. Therefore, before creating a migration task in a version of Tiered Storage Manager earlier than 5.9, check when storage domain information was last updated in the **Detailed Information** tab of the *storage-domain-name* subwindow, and then refresh the storage domain if necessary.

In these versions, if you change the LDEV configuration, hardware configuration, or other information in Device Manager but do not update the change to Tiered Storage Manager, volume information will no longer be displayed correctly. If you then create and execute a migration task without refreshing, data might be migrated to an unexpected location.

The number of candidate volumes to be displayed

In versions of the Tiered Storage Manager server earlier than 5.6, the information of all candidate volumes in the migration target storage tier was displayed when creating a migration plan. For this reason, displaying the information took a long time, and a memory shortage sometimes occurred. In Tiered Storage Manager 5.7 or later, however, only enough candidate volumes to equal the number of migration source volumes are displayed.

To display all candidate volumes, change the setting of the property `server.migrationPlan.candidateVolumeCountLimit` in the file `server.properties`.

Change in the default setting for deleting data

In versions of the Tiered Storage Manager server earlier than 5.6, migration source volumes were automatically deleted after migration by default. In Tiered Storage Manager 5.7 or later, however, to protect important data by avoiding accidental erasure, migration source volumes are not deleted by default.

To avoid information leakage, we recommend that you shred the data of migration source volumes after migration.

Display of task owner names

In Tiered Storage Manager 5.7 or later, the name of a task owner, which is a task attribute, is the user ID specified during user registration, not the user ID entered during login.

User IDs entered during login are not case sensitive. For this reason, even if the same user creates multiple tasks, the upper-case and lower-case characters in the task owner name for a task created using Tiered Storage Manager earlier than 5.7 might differ from the owner name for a task created using Tiered Storage Manager 5.7 or later.

Glossary

This glossary defines the special terms used in this document. Click the desired letter below to display the glossary entries that start with that letter.

C

Can Migrate

One of the attributes that is set when creating a migration group. The Can Migrate attribute indicates whether the migration group can be selected as a migration source. If **Yes** is selected, the migration group can be selected as a migration source. If **No** is selected, the migration group cannot be selected as a migration source. On the Web client, these migration groups are not displayed as migration source candidates.

CLPR (Cache Logical Partition)

A Universal Storage Platform V/VM and Hitachi USP function that manages the cache of a storage subsystem by logically partitioning it. Migration within the same storage domain is possible even between volumes partitioned using the CLPR functionality.

D

data erasure

A function to erase the data on the migration source volumes after completing a migration.

Device Manager

Device Manager is a prerequisite program product of Tiered Storage Manager. Device Manager can manage a system consisting of Hitachi storage subsystems with a unified view.

DEVN (device number)

A number that is used to identify a volume on a mainframe host. This number is different from the device number of Tiered Storage Manager. To display a DEVN in Tiered Storage Manager, you need to install Mainframe Agent on the host and register the host in Device Manager.

domain controller

The storage subsystem used to control a storage domain. Universal Storage Platform V/VM and Hitachi USP can be used as domain controllers.

DP pool

A storage area for storing write data from hosts to a DP volume. The Dynamic Provisioning function of Universal Storage Platform V/VM provides this area by using at least one real volume.

DP pool volume

A real volume that makes up a DP volume.

DP volume

A virtual volume that can be used for the Dynamic Provisioning function of Universal Storage Platform V/VM.

DP volume group

A group of virtual volumes created using Open Volume Management.

Dynamic Provisioning

A function of Universal Storage Platform V/VM. This function allows a host to monitor available pool capacity, and, in accordance with a write request from the host, provides a volume (DP volume) to which necessary physical disks are allocated.

E

emulation type

An LDEV attribute used when an LDEV is created in an array group of the Universal Storage Platform V/VM, Hitachi USP, or Lightning 9900 storage subsystems. The LDEV size is determined from this attribute.

EVA

HP StorageWorks Enterprise Virtual Array

event notification function

A function that uses email to report the occurrence of the following events, which are not directly the result of user actions:

Completion of migration tasks, Completion of shredding tasks, Completion of locking tasks, Completion of unlocking tasks, Completion of volume creation tasks, Completion of external connection setup tasks, Expiration of a volume lock period, Expiration of a specified period

external storage subsystem

A storage subsystem connected to the domain controller by the external storage connection function provided by Universal Volume Manager. Tiered Storage Manager registers an external storage subsystem as belonging to the same storage domain as the internal storage subsystem (Universal Storage Platform V/VM or Hitachi USP) it is connected to.

external volume

A volume on an external storage subsystem.

H

Hitachi Storage Command Suite Common Component

A component that provides functions shared by Hitachi Storage Command Suite products. It is installed as a part of Tiered Storage Manager, and provides functions such as login, integrated log output, and Web services.

Hitachi Storage Command Suite products

A general term for Hitachi storage-related products. Versions earlier than 6.0 used the brand name HiCommand. Since version 6.0, however, the brand name is Hitachi Storage Command Suite.

I

internal volume

A volume in Universal Storage Platform V/VM or Hitachi USP. Specifically, this term is used to distinguish such volumes from external volumes.

L

LDEV (logical device)

A logical partition of storage areas in an array group of a storage subsystem. An LDEV is represented by a combination of the storage subsystem's product name, storage subsystem's serial number, and internal LU. Also called a *logical device*.

logical DKC

A term used in volume management when managing volumes in Universal Storage Platform V/VM in 64 KB LDEV units. In Universal

Storage Platform V/VM, migration is possible only if the source volume and the target volume have the same logical DKC. This requirement results in storage domains being created in units of logical DKCs.

logical DKC number

The number assigned to identify a logical DKC.

logical DKC serial number

The logical device identification number assigned to a logical DKC. This number is displayed as *logical-DKC-number* + *serial-number*.

logical group

A Device Manager function that allows users to manage the paths to the volumes assigned to a host by dividing the paths into hierarchical groups.

LU (logical unit)

A logical disk disclosed to a host as a SCSI LDEV on a storage subsystem port.

LUSE (LU size expansion)

Universal Storage Platform V/VM and Thunder/Lightning functions that increase volume capacity by virtually integrating LUs. Device Manager can set up LUSE volumes.

M

management server

A computer on which the server program of Tiered Storage Manager runs.

message log

A history of user-executed processes and processing results. Information output to the message log is more detailed than information output to the Hitachi Storage Command Suite common log.

migration

Relocation of application data from the current storage location to another storage location. This change in data storage location does not affect applications.

migration plan

An object where a migration group that serves as the migration source and a storage tier that serves as the migration destination are specified.

You can create or edit a migration plan in the migration task management window on a Web client, or in a file format on a CLI client.

N

normal volume

In Tiered Storage Manager, the normal volume indicates a volume other than a virtual volume or pool volume.

P

pool

A storage area for virtual volumes.

pool volume

Real volumes that make up a pool.

port/host storage domain

Port indicates the channel adapter port or port controller port. Host storage domain indicates a group of WWNs that are allowed to issue I/O to the specified volumes. You can specify multiple volumes. A host storage domain is managed by the LUN Manager of the storage subsystem and can also be called a host group.

port/host storage domain is a value specified by connecting a port name and host storage domain name with a forward slash (/).

property file

A file used for settings such as server operation and default options for CLI execution. Hitachi Tiered Storage Manager has the following server property files:

`server.properties` file

`database.properties` file

`devicemanager.properties` file

`logger.properties` file

`tuningmanager.properties` file

There are two CLI property files:

`htsmcli.properties` file

`htsmclienv.properties` file

PSA (partitioned storage administrator)

An administrator for partitioned volumes that use SLPR functionality of Universal Storage Platform V/VM or Hitachi USP. A domain controller for which a PSA manages operations cannot be registered as a storage

domain. Therefore, Tiered Storage Manager cannot be used. Operation is not guaranteed even when an SA (overall storage administrator) registers a storage domain and the user is changed to a PSA.

PSUE (pair suspended-error)

A status that indicates that an error has occurred for a copy pair in a mainframe host.

PSUS (pair suspended-split)

A status that indicates that an update of a P-VOL has not been applied to an S-VOL in a mainframe host.

R

real volume

A volume in which all the areas of a specified capacity were already allocated when the volume was created in Universal Storage Platform V/VM. In this manual, the real volume indicates a volume other than a virtual volume.

S

serial number

The device identification number assigned to the storage subsystem.

SLPR functionality

A Universal Storage Platform V/VM and Hitachi USP function that manages the volumes of all storage subsystems by logically partitioning them into SLPRs. If volumes that are logically partitioned are not managed by a PSA (partitioned storage administrator), migration is possible even between volumes that have different partitions. A PSA cannot operate Tiered Storage Manager.

SYSPLEXID

A name assigned by a user to a mainframe host to identify a volume on the host. To display the SYSPLEXID in Tiered Storage Manager, you need to install Mainframe Agent on the host and register the host into Device Manager.

U

Universal Volume Manager

A program that enables storage subsystems to establish external connections with Universal Storage Platform V/VM or Hitachi USP. For details on storage subsystems that support external connections to

Universal Storage Platform V/VM and Hitachi USP, see the *Universal Volume Manager User's Guide*.

V

virtual volume

A generic term for one of the two types of virtual volumes used in storage functions:

Copy-On-Write Snapshot V-VOL

DP volume

volume

A storage area in which application data is stored. A logical unit (LU) is also a type of volume.

Acronyms and Abbreviations

There are numerous acronyms and abbreviations specific to storage technologies. Following is a list of acronyms and abbreviations that are used in describing concepts and procedures common to Tiered Storage Management.

B

BD

Basic Disk

C

CLI

command line interface

CLPR

cache logical partition

CVS

custom volume size

D

DEVN

device number

DKC

disk controller

DLCM

Data Lifecycle Management

DP

Dynamic Provisioning

F

FC

A fiber channel that functions as a serial I/O interconnect supporting multiple protocols.

G

GB

gigabyte(s)

H

HD-FC

high density fiber channel

K

KB

kilobyte(s)

L

LDAP

Lightweight Directory Access Protocol

LAN

local-area network

LDEV

logical device

P

P-VOL

primary volume

PSUE

pair suspended-error

PSUS

pair suspended-split

R

RADIUS

Remote Authentication Dial In User Service

S

SAS

Serial attached SCSI that provides interfaces for attaching buses and controllers that connect to other devices, including disk and tape drives

SATA

serial advanced technology attachment

S-VOL

secondary volume

SLO

service level objective

SLPR

storage management logical partition

SNIA

Storage Networking Industry Association

SNMP

Simple Network Management Protocol

SVP

service processor

T

TCP/IP

transmission control protocol/internet protocol

V

VMA

volume management area

VOLSER

volume serial number

W

WWN

Worldwide Name

Index

A

adding
  user 3–8
application area 3–3
application bar area 3–3
authentication
  changing user authentication method 3–15

C

canceling
  tasks 3–42
changing
  password 3–10
checking property information
  migration group 3–34
  storage domain 3–20
  storage tier 3–26
creating
  migration group by using Create MG dialog box 3–31
  storage tier by specifying volume filter conditions 3–24
  storage tier from Create Tier dialog box 3–22

D

Dashboard 3–2
deleting
  migration group 3–35
  storage domain 3–21
  storage tier 3–28
  tasks 3–42
  user 3–14
dialog box 3–3

E

editing attribute information
  migration group 3–34
  storage tier 3–27
  tasks 3–43
editing properties
  storage domain 3–20
excluding
  volumes 3–38
Explorer 3–2
external authentication server 3–15

F

filter
  migration group based on 3–32

G

global task area 3–2

L

LDAP 3–15
lock status 3–13

M

message logs
  file name 4–9
  output destination 4–9
migration
  excluding source volumes 3–38
  performing 1–5
  planning 1–3
  preparing 1–4
  purpose 1–2
  steps 1–3
migration group
  checking property information 3–34
  creating by filtering volumes 3–32
  creating, using the Create MG dialog box 3–31
  deleting 3–35
  editing attribute information 3–34
migration volume
  adding volumes 3–32

N

navigation area 3–3

P

password
changing 3–10
conditions 3–12
performing
migration 1–5
permissions
user 3–9
planning
migration 1–3
preparing
migration 1–4
profile
user 3–10

R

RADIUS 3–15
Reference Information A–1
refresh statuses 2–26, 2–29

S

searching
tasks 3–41
storage domain
checking property information 3–20
deleting 3–21
storage tier
checking property information 3–26
creating by specifying volume filter conditions 3–24
creating, from Create Tier dialog box 3–22
deleting 3–28
editing attribute information 3–27
relationship between storage tiers and volumes 2–3
subwindow 3–3

T

tasks
canceling 3–42
checking status 3–41
deleting 3–42
editing attribute information 3–43
executing 3–40
searching for 3–41
stopping execution 3–41
trace logs
file name 4–10
output destination 4–10

U

user
adding 3–8
automatic locking 3–12
deleting 3–14
lock status 3–13
permissions 3–9
profile 3–10
user authentication method
changing 3–15
User Management 3–6

V

volume
adding 3–32
volumes
grouping 2–7, 2–10

W

warning banner 3–19

Hitachi Data Systems

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com

Hitachi Data Systems

MK-94HC090-12