HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™ Lightning 9900™

Dynamic Optimizer User's Guide

[Hitachi Internal Hierarchical Storage Manager (HIHSM)]

Document Revision Level

Revision	Date	Description
MK-90RD007-P	May 2000	Preliminary Release
MK-90RD007-P1	June 2000	Rev. 1 of Prelim. Release, supersedes and replaces MK-90RD007-P
MK-90RD007-0	July 2000	Initial Release
MK-90RD007-1	November 2000	Revision 1, supersedes and replaces MK-90RD007-0

Source Document Revision Level

The following source documents were used to produce this 9900 user guide:

- *DKC310 Disk Subsystem, Hitachi Internal Hierarchical Storage Manager Reference Manual*, revision 4.
- MK-90RD007-S1 (STR RAID400 document).
- ECN TW-90RD003-0.

Changes in this Revision

This revision includes the following changes and additions:

- The HIHSM graphical user interface has been improved (see Figures in Chapter 4).
- Added the new GraphTool software product (replaces Microsoft® Excel usage): sections 1.2, 2.4, 2.8, Chapter 4, and new Chapter 5 (old section 4.9).
- Corrections: “disaster recovery and regeneration” to “data recovery and regeneration”.
- Added a note on the license key codes (see section 3.3).

Preface

The *Hitachi Lightning 9900™ Dynamic Optimizer (HIHSM) User's Guide* describes and provides instructions for using the Hitachi Internal Hierarchical Storage Manager (HIHSM) remote console software to configure and perform Dynamic Optimizer operations on the Hitachi Lightning 9900™ RAID storage subsystem. This user's guide assumes that:

- the user has a background in data processing and understands direct-access storage device (DASD) subsystems and their basic functions,
- the user has read and understands both the *Hitachi Lightning 9900™ User and Reference Guide* and the *9900 Remote Console User's Guide*, and
- the user is familiar with the Microsoft® Windows® 95/98/NT operating systems (e.g., opening, closing, minimizing, and restoring windows; using the keyboard and mouse to navigate on screen and select objects).

Note: The term "9900" refers to the entire Hitachi Lightning 9900™ subsystem family, unless otherwise noted. Please refer to the *Hitachi Lightning 9900™ User and Reference Guide* (MK-90RD008) for further information on the 9900 disk array subsystems.

Note: The use of Dynamic Optimizer, the HIHSM software, and all other Hitachi Data Systems products is governed by the terms of your license agreement(s) with Hitachi Data Systems.

Contents

List of Figures

List of Tables

Chapter 1 Overview of Dynamic Optimizer and HIHSM

1.1 Dynamic Optimizer

The Dynamic Optimizer feature of the Hitachi Lightning 9900™ RAID subsystem enables you to optimize your data storage and retrieval on the 9900 subsystem. Dynamic Optimizer analyzes detailed information on the usage of 9900 subsystem resources and tunes the 9900 automatically by migrating logical volumes within the subsystem according to detailed user-specified parameters. Dynamic Optimizer tuning operations can be used to resolve bottlenecks of activity and optimize volume allocation. Dynamic Optimizer operations are completely nondisruptive; the data being migrated can remain online to all hosts for read and write I/O operations throughout the entire volume migration process. Dynamic Optimizer also supports manual volume migration operations and can estimate performance improvements prior to migration to assist you in tuning the 9900 subsystem for your operational environment.

Dynamic Optimizer provides the following major benefits for the user:

- **Load balancing of subsystem resources.** Balancing resource utilization can significantly improve 9900 subsystem performance. The data provided by Dynamic Optimizer enables you to optimize several areas of performance, including front-end and back-end processor usage as well as the allocation of logical volumes to physical drives and RAID level.
- **Optimizing disk drive access patterns.** Dynamic Optimizer collects and analyzes detailed information on disk drive access patterns and can migrate volumes to optimize host access to the data stored on the 9900 subsystem. For example, RAID-1 technology may provide better performance than RAID-5 under certain operational conditions, and one disk drive type may provide better performance than another for certain types of access. The Dynamic Optimizer feature enables you to fine-tune the logical volume allocation of the 9900 subsystem to optimize host access to data.

The 9900 subsystem supports both RAID-1 and RAID-5 technologies as well as an intermix of RAID-1 and RAID-5 parity groups. The 9900 subsystem also supports several types of hard disk drives (HDDs) and allows an intermix of HDD types within each array domain to provide maximum flexibility in configuration. Dynamic Optimizer operations take into account the RAID level and physical HDD performance of each parity group, enabling reallocation of logical volumes and optimization with respect to both RAID level and HDD type. The proper combination of RAID level and HDD type for the logical volumes can significantly improve 9900 subsystem performance for the user's environment.

In addition to RAID level and HDD type, Dynamic Optimizer bases its 9900 tuning plans on logical device usage, back-end path usage, and disk drive access patterns. Dynamic Optimizer also applies detailed user-specified criteria to plan and perform 9900 performance tuning. For example, Dynamic Optimizer allows you to specify the maximum disk utilization for each class of installed HDD, select the range of 9900 performance data to be analyzed, exclude specific parity groups from internal migration operations, schedule the automatic migration operations, and limit the time and performance impact of the migration operations.

1.2 HIHSM Remote Console Software

The licensed Hitachi Internal Hierarchical Storage Manager (HIHSM) software for the 9900 Remote Console PC allows you to configure and perform Dynamic Optimizer operations on the 9900 subsystem. The 9900 Remote Console PC is attached to and communicates directly with the 9900 subsystems via the 9900-internal LAN. For further information on the 9900 Remote Console PC, please refer to the *9900 Remote Console User's Guide*.

The HIHSM software functions include:

- Monitoring. The HIHSM monitor function monitors, collects, and displays usage statistics for various 9900 subsystem resources, including parity groups, logical volumes, and front-end and back-end microprocessors.
- Estimating. The HIHSM estimate function calculates the expected usage of the source and target parity groups after a proposed volume migration. The estimate function uses the subsystem usage data collected by the monitor function in addition to the HDD performance characteristics and RAID level to estimate expected usage rates. Both the auto migration and manual migration functions use these estimated values to verify proposed migrations.
- Auto migration. The HIHSM auto migration function analyzes monitor data and creates and executes migration plans automatically based on user-specified parameters. HIHSM allows you to schedule the auto migration operations, select the monitor data to be analyzed, and specify the maximum time and maximum performance impact of the auto migration operations. The HIHSM software displays a detailed log of auto migration activities.
- Manual migration. The HIHSM manual migration function allows you to select and migrate logical volumes within the 9900 subsystem. The manual migration function displays the estimated results of your proposed migration operations, enabling you to verify expected performance improvements prior to actual migration. The HIHSM software displays a detailed log of manual migration operations.
- Reserving volumes. The HIHSM reserve volume function allows you to select and reserve the desired target volumes for auto and manual migration operations. The reserve volume function provides an additional level of user control over Dynamic Optimizer operations.
- Data graphing. The HIHSM GraphTool data graphing software allows you to extract the HIHSM monitor data from the export file and graph the data in several data plotting formats. GraphTool allows you to perform your own analysis of the 9900 hardware usage data to identify workload characteristics, peaks and trends in usage, and bottlenecks of front-end and back-end activity.

Note: The use of the HIHSM software and all Hitachi Data Systems products is governed by the terms of your license agreement(s) with Hitachi Data Systems.

Chapter 2 Overview of Dynamic Optimizer Operations

Dynamic Optimizer tuning operations are configured and performed using the licensed HIHSM software on the 9900 Remote Console PC. The HIHSM software allows you to monitor the usage of 9900 subsystem resources and plan and perform subsystem-internal volume migration operations to optimize subsystem performance. HIHSM provides both automatic and manual volume migration capabilities. For automatic migration, the volumes to be migrated are selected according to user-specified parameters, and the volume migration operations are performed automatically according to the user-specified schedule. For manual migration, HIHSM displays the expected results of user-proposed migration operations, and performs user-specified volume migration operations immediately upon request.

2.1 Components

The 9900 subsystem components which are involved in Dynamic Optimizer operations are:

- HIHSM remote console software (see section 2.1.1),
- Channel processors (CHPs) (see section 2.1.2),
- Disk processors (DKPs) (see section 2.1.3),
- Data recovery and regeneration processors (DRRs) (see section 2.1.4),
- Access paths (adapter-CSW, CSW-cache, and adapter-SM) (see section 2.1.5),
- Logical volumes (see section 2.1.6),
- Parity groups (see section 2.1.7), and
- Hard disk drives (HDDs) (see section 2.1.8).

2.1.1 HIHSM Remote Console Software

The HIHSM remote console software enables you to perform Dynamic Optimizer operations on the 9900 subsystems attached to the Remote Console PC. The 9900 Remote Console PC is attached to the 9900 subsystem(s) via the 9900-internal LAN and communicates directly with the 9900 service processor (SVP). For further information on the 9900 Remote Console PC, please refer to the *9900 Remote Console User's Guide*.

Note: The HIHSM software will not function on a 9900 subsystem which does not have the Dynamic Optimizer feature installed and enabled.

2.1.2 Channel Processors (CHPs)

The CHPs are the microprocessors on the client-host interface processor (CHIP) boards which process the channel commands from the hosts and manage host access to cache. The 9900 can be configured to provide up to 32 CHPs. The HIHSM software monitors and displays the usage of all CHPs.

2.1.3 Disk Processors (DKPs)

The DKPs are the microprocessors on the array control processor (ACP) boards which control the transfer of data between the disk drives and cache. The 9900 subsystem can be configured to provide up to 32 DKPs. The HIHSM software monitors and displays the usage of all DKPs.

2.1.4 Data Recovery and Regeneration Processors (DRRs)

The DRRs are the microprocessors, also located on the ACPs, which are used to generate parity data for RAID-5 parity groups. The DRR uses "old data + new data +old parity" to generate new parity. The HIHSM software monitors and displays the usage of all DRRs.

2.1.5 Access Paths (Adapter-CSW, CSW-Cache, and Adapter-SM)

The 9900 subsystem utilizes cache path switches (CSWs) having a cache memory arbiter (CARB). Each CHP, DKP, and cache element is attached to the CSWs, and the CSWs are used simultaneously for both data and command functions. This architecture is called the hierarchical star net interface. The hierarchical star net (HSN) interface, unlike the bus interface on the 7700E subsystem, provides overall access reliability in case of a failure anywhere in the subsystem. Because of the new HSN architecture, the failure of one or two CHPs/DKPs will not bring the entire subsystem down. All data stored on the 9900 subsystem is moved to and from cache via the CSWs; there is no direct movement of data between CHPs and DKPs. The HIHSM software monitors and displays the usage of all 9900 access paths between CHP/DKP and CSW, between CSW and cache, and between CHP/DKP and shared memory (SM).

2.1.6 Logical Volumes

The 9900 subsystem supports a maximum of 4,096 logical volumes (256 volumes per logical control unit (CU) image). The number of logical volumes per parity group depends on the HDD type, RAID level, and device emulation type (e.g., 3390-3R, OPEN-9) of the parity group. Each logical volume is uniquely identified by its three-digit logical device (LDEV) ID, which consists of the logical CU image number (0, 1, 2…F) and device number (00-FF hexadecimal). The HIHSM software monitors and displays the usage of all logical volumes.

2.1.7 Parity Groups

A parity group, also called an array group, is a group of hard disk drives (HDDs) which forms the basic unit of storage for the 9900 subsystem. All HDDs in a parity group must have the same physical capacity. Each parity group is attached to both ACPs of an ACP pair, enabling all HDDs in the group to be accessed simultaneously by the ACP pair. (Each ACP pair contains four DKPs.) The 9900 supports both RAID-1 and RAID-5 parity groups. A RAID-1 parity group consists of two pairs of HDDs in a mirrored configuration. A RAID-5 parity group consists of four HDDs. Figure 2.1 shows a RAID-5 parity group consisting of four 18-GB HDDs and the RAID data striping across the HDDs. The HIHSM software monitors and displays the usage of all parity groups.

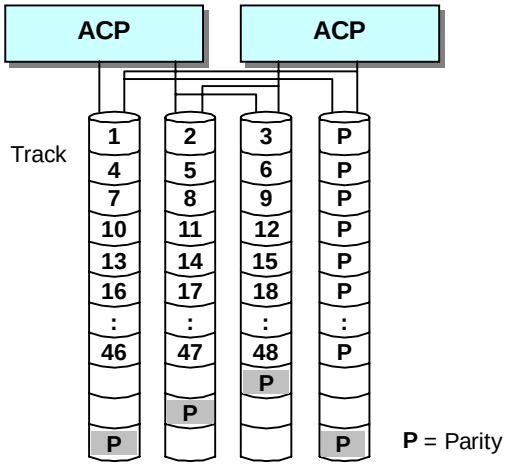

Figure 2.1 RAID-5 Parity Group and Data Striping

2.1.8 Hard Disk Drives (HDDs)

The 9900 HDDs are designed for the large enterprise-systems market and meet the most demanding specifications. Disk drives of varying capacities can be attached to the same ACP pair, but all drives in a parity group must have the same capacity. The 9900 automatically detects and corrects disk errors and invokes dynamic sparing if needed. For both RAID-1 and RAID-5 parity groups, a spare disk can back up any other disk of the same capacity anywhere in the subsystem.

The 9900 HDD types are ordered according to level of performance, as shown in Table 2.1. For each 9900 subsystem, the HIHSM software classifies the installed HDD with the highest performance as Class A, the installed HDD with the next highest performance as Class B, and so on, with one class for each installed HDD type. Figure 2.2 shows two examples of HIHSM class assignments for installed HDDs.

Table 2.1 HDD Types versus Performance

Order	HDD Type	Capacity	
0	DKR2B	18 GB	*High Performance*
1	DKR1C	72 GB	
2	DKR1B	47 GB	*Low Performance*

Note: For the latest information on available HDD types, please contact your Hitachi Data Systems account team.

Example 1

Installed HDD	Order	HIHSM Class
DKR2B	0	Class A
DKR1C	1	Class B
DKR1B	2	Class C

Example 2

Installed HDD	Order	HIHSM Class
DKR2B	0	Class A
DKR1B	2	Class B

Figure 2.2 Examples of Class Assignments for Installed HDD Types

2.2 Volume Migration Operation

The HIHSM volume migration operation consists of two steps:

- Copy the data on the HIHSM source volume to the HIHSM target volume (see Figure 2.3).
- Transfer host access to the target volume to complete the migration (see Figure 2.4).

The HIHSM source volume can be online to all hosts during the migration operation, and the HIHSM target volume is reserved prior to migration to prevent host access during migration. The HIHSM copy operation copies the entire contents of the source volume to the target volume cylinder by cylinder (24 tracks at a time, not including the diagnostic and unassigned alternate tracks). If the source volume is updated by write I/Os during the HIHSM copy operation, the 9900 keeps track of the updates on a cylinder map of the source volume, and performs additional copy operations after the HIHSM copy operation is complete to duplicate the updates at the target volume. When the volumes are fully synchronized (no differential data on the source volume), the 9900 subsystem completes the migration operation by redirecting host access to the target volume.

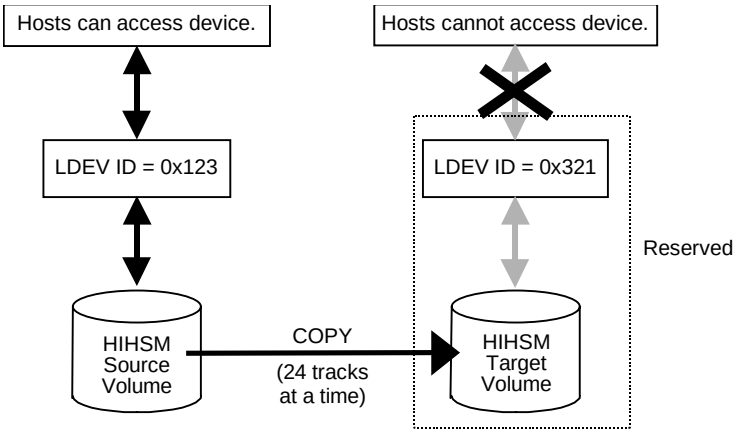

Figure 2.3 During a HIHSM Volume Migration Operation

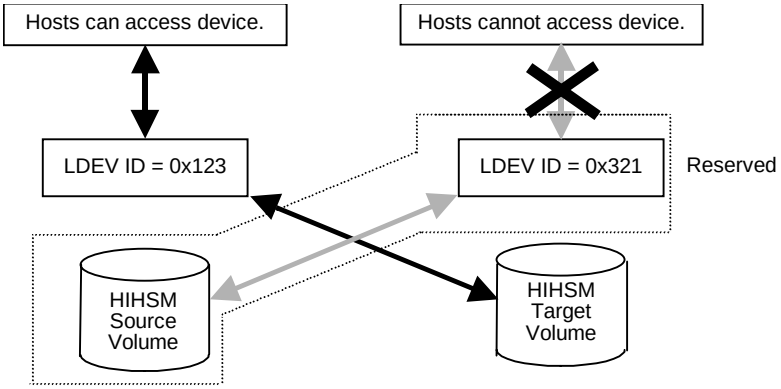

Figure 2.4 After a HIHSM Volume Migration Operation

2.3 Requirements and Restrictions

The requirements and restrictions for performing Dynamic Optimizer operations are:

- **Logical volumes.** The Dynamic Optimizer source and target volumes must be specified by LDEV ID, not VOLSER or port/TID/LUN. The Dynamic Optimizer source and target volumes must be within the same 9900 subsystem and have the same device emulation type and same size. Dynamic Optimizer does not support the following logical volumes:
  - Multiplatform volumes (e.g., 3390-3A/B/C, 3380-KA/B/C).
  - Volumes which are set as command devices (devices reserved for use by the host).
  - Volumes which are assigned to Hitachi ShadowImage (HMRCF/HOMRCF), Hitachi Remote Copy (HRC/HORC), or Hitachi Online Data Migration (HODM) pairs. *Caution:* If you delete an HRC/HORC pair from the RCU, or split a ShadowImage pair, the HIHSM software will display this simplex volume as available for migration.
  - Volumes which are reserved for ShadowImage operations.
  - Volumes which have FlashAccess (also called DCR) data stored in cache.
  - Volumes which are in an abnormal or inaccessible condition (e.g., pinned track, fenced).
- **Target volumes.** Target volumes must be reserved prior to migration. The HIHSM remote console software allows you to reserve volumes as HIHSM target volumes.
  - Hosts cannot access HIHSM-reserved volumes.
  - Logical Unit Size Expansion (LUSE) volumes cannot be reserved as target volumes.
  - The target volume must be the same size and emulation type as the source volume.
- **LUSE source volumes.** If you want to specify a LUSE source volume for migration, you can specify individual LDEVs within the LUSE volume (e.g., the LDEVs with high access rates), and HIHSM will migrate only the specified LDEVs of the LUSE volume. If desired, you can also specify all LDEVs of the LUSE volume to relocate the entire volume. In this case, you must make sure that the required number of reserved target LDEVs is available.
- **Number of volumes:**
  - The maximum number of HIHSM copy operations that can be requested at the same time is 36.
  - The maximum number of HIHSM copy operations plus ShadowImage pairs (HMRCF/HOMRCF) is 1024.
- **Auto migration plan.** HIHSM will **not** be able to make an auto migration plan when:
  - HIHSM cannot estimate the usage rate for all migrated parity groups. The estimate function calculates the expected results of a proposed volume migration based on the 9900 monitor data. If any condition prevents HIHSM from estimating an expected usage rate (e.g., invalid monitor data), HIHSM will not be able to make the plan.
  - The estimated usage rate of any migrated parity group is over the maximum rate. If any estimated usage rate is over the limit, HIHSM will not be able to make the plan.
- **9900 maintenance.** Do not perform HIHSM migration operations during subsystem maintenance activities (e.g., cache or drive installation, replacement, or deinstallation).

2.4 Monitoring Subsystem Usage

HIHSM monitoring provides the detailed subsystem usage information which the auto migration function analyzes to determine whether any resources are overloaded or unbalanced. The HIHSM remote console software allows you to configure, start, and stop HIHSM monitoring and select the monitor data term (range of data) to be displayed on screen and used in analyses. The HIHSM GraphTool software allows you to extract the monitor data from the export file and graph this statistical information in several data plotting formats.

HIHSM monitor data is collected twice per day automatically as requested by the user (four data samples per hour). The 9900 SVP stores the last three months (90 days) of HIHSM monitor data on its hard disk drive. When the HIHSM monitor data is more than three months old, the oldest HIHSM data is overwritten by the new HIHSM data, ensuring that current HIHSM monitor data is always available. ***Note:*** The user is responsible for exporting HIHSM monitor data as needed to save old HIHSM data before it is overwritten.

The HIHSM monitor function collects and displays the following 9900 subsystem usage data:

- **Parity group usage**: The HIHSM Main panel (see section 4.5.1) displays the average and maximum usage of each parity group. The parity group usage value is the sum of the logical volume usage values for all volumes in the parity group.

  If HIHSM monitoring shows overall high parity group usage, you should consider installing additional HDDs and then using HIHSM to migrate the high-usage volumes to the new parity groups. If HIHSM monitoring shows that parity group usage is not balanced, you can use HIHSM to migrate volumes from high-usage parity groups to low-usage parity groups.

- **Logical volume usage**: The Volume List panel (see section 4.5.2) displays the average and maximum usage, including sequential and random access, of each logical volume in a parity group. The logical volume usage is the time in use (sequential and random access) of the physical drives of each logical volume, averaged by the number of physical drives in the parity group.

  If HIHSM monitoring shows overall high logical volume usage, you should consider installing additional hardware (e.g., HDDs, ACP pairs, cache). If HIHSM monitoring shows that volume usage is not balanced, you can use HIHSM to migrate high-usage volumes to higher HDD classes and/or to lower-usage parity groups.

- **CHP usage**: The Usage panel (see section 4.5.3) displays the average and maximum usage of each CHP.

  If HIHSM monitoring shows overall high CHP usage, you should consider installing additional CHPs (CHIP pairs). If HIHSM monitoring shows that CHP usage is not balanced, you should consider reconnecting the channel paths connected to the overloaded CHPs to a different CHIP pair which contains lower-usage CHPs. HIHSM migration operations do not address and will not improve CHP usage.

- **DKP usage**: The Usage panel displays the average and maximum usage of each DKP.

  If HIHSM monitoring shows overall high DKP usage, you should consider installing additional HDDs and/or ACP pairs, and then using HIHSM to migrate the high-write-usage volumes (especially sequential writes) to the new parity groups. If HIHSM monitoring shows that DKP usage is not balanced, you can use HIHSM to migrate logical volumes from high-usage ACP pairs to low-usage ACP pairs. (Each ACP pair contains four DKPs.)

  Note: HIHSM cannot estimate DKP usage. Only use HIHSM migration for obvious cases of high or unbalanced DKP usage. HIHSM migration may not provide any performance improvement for cases in which DKP usage values vary only slightly, or for cases in which overall DRR usage values are relatively high.

- **DRR usage**: The Usage panel displays the average and maximum usage of each DRR.

  If HIHSM monitoring shows overall high DRR usage, this can indicate high write penalty. Please consult your Hitachi Data Systems representative about high write penalty conditions. If HIHSM monitoring shows that DRR usage is not balanced, you should consider relocating volumes using HIHSM to balance DRR usage within the subsystem.

- **Access path usage**: The Usage panel displays the average and maximum usage of each access path (adapter-CSW, CSW-cache, and adapter-SM) for the specified monitor term.

  HIHSM monitoring displays access path usage information, but since the access paths are used simultaneously by all 9900 CHPs and DKPs, HIHSM volume migration operations do not address and will not improve access path usage. Please contact your Hitachi Data Systems representative if you have questions about your 9900 access path usage data.

- **Write pending rate**: The Usage panel displays the average and maximum write-pending rate for the specified monitor term.

Note: HIHSM migration should be performed only when you can expect a large improvement in subsystem performance. HIHSM migration may not provide significant improvement for cases in which parity group or volume usage varies only slightly, or for cases in which overall DKP or DRR usage is relatively high. Also keep in mind that subsystem tuning operations can improve performance in one area while at the same time decreasing performance in another. For example, let's say that parity groups A and B have average usage values of 20% and 90%, respectively. HIHSM estimates that if one logical volume is migrated from parity group B to parity group A, the usage values will become 55% and 55%. If you perform this migration operation, the I/O response time for parity group B will probably decrease, and the I/O response time for parity group A may increase, while the overall throughput may increase or decrease.

Caution: When an error condition exists in the 9900 subsystem, resource usage can increase or become unbalanced. Do not use 9900 monitor data collected during an error condition as the basis for planning HIHSM migration operations.

2.5 Estimating Usage Rates

The HIHSM software automatically estimates the expected usage rates for the source and target parity groups for each proposed migration operation. The estimate function calculates the expected results of a proposed volume migration based on the 9900 monitor data collected during the user-specified monitor data term. The HIHSM estimate function takes the RAID level and HDD type into account when estimating expected usage rates (RAID-1 and RAID-5 have different access characteristics).

The HIHSM estimate function is a critical component of the HIHSM auto migration function. The estimate function calculates the expected parity group usage for both the source and target parity groups, and the auto migration function uses these estimated values to verify proposed migrations. If any condition prevents HIHSM from estimating an expected usage rate (e.g., invalid monitor data), the HIHSM auto migration function will not be able to make the plan.

For manual migration operations, when you select the source volume for a manual migration, HIHSM displays the expected usage rates for the source parity group, so that you can see the predicted effect of migrating the selected volume out of its group. When you select the target volume for a manual migration, HIHSM displays the expected usage rates for the target parity group, so that you can see the predicted effect of adding the selected source volume to the selected parity group.

Note: HIHSM does not estimate processor or access path usage. HIHSM migration operations can be used to improve DKP and DRR usage, but do not address CHP or access path usage. Please refer to section 2.4 for further information on interpreting the CHP, DKP, DRR, and access path usage data.

2.6 Automatic Migration Operations

The HIHSM auto migration function performs volume migration automatically based on user-specified parameters and is intended to function as the primary tuning method for the 9900 subsystem. HIHSM gives you full control over the auto migration function, enabling you to specify the parameters and schedule for all auto migration operations. HIHSM also enables you to limit the impact of auto migration on subsystem performance by specifying both usage and time limits for the auto migration operations.

HIHSM auto migration operations are based on the following three major parameters:

- **Hierarchy of parity groups.** HIHSM arranges the parity groups in the 9900 subsystem in a hierarchy, or ordering, based on HDD type. HIHSM assigns each parity group to a class based on the performance of its HDD type, and the classes are ordered from highest performance drive (class A) to lowest performance drive (class B-C or higher depending on the number of HDD types available). The HIHSM auto migration function uses this hierarchy to identify target volumes for auto migration operations.

  The HIHSM auto migration function also allows users to specify "fixed" parity groups which are to be excluded from auto migration operations. HIHSM will not migrate any volumes into or out of fixed parity groups.

- **Maximum disk usage.** The HIHSM auto migration function allows you to specify the maximum disk usage rate for each HDD class in the 9900 subsystem, and then uses these limits to identify source volumes for auto migration operations. The HIHSM auto migration function uses only the disk usage data to identify source volumes. Processor and access path usage data are not used by the auto migration function. For information on evaluating processor and access path usage data, please refer to section 2.4.

  You must identify and specify the appropriate disk usage limits for your operational environment. When you use the same maximum disk usage rate for all HDD classes, the performance of the HDD is the only factor used in determining auto migration plans. When you specify different usage limits for the HDD classes, you can bias the auto migration function to favor (or avoid) certain HDD types. The migration of high-usage volumes to higher HDD classes is expected to significantly improve host access to the volumes, which in itself can also have a large effect on subsystem performance.

  The HIHSM auto migration function also allows you to specify the maximum disk usage rate during auto migration operations, so that you can control the impact of HIHSM copy operations on subsystem performance. If the usage of the source or target parity group exceeds the specified limit during migration, HIHSM aborts that auto migration operation.

- **Monitor data term.** The HIHSM auto migration function uses the disk usage data collected during the user-specified monitor data term to make auto migration plans. ***Note:*** HIHSM will not use any monitor data collected before the last volume migration. This eliminates the possibility of auto migration plans being based on invalid data.

Note: Do not perform manual migration operations while the auto migration function is active. Always turn off the auto migration function and delete all existing auto migration plans before performing manual migration.

Storage management by maximum disk usage. The HIHSM auto migration function allows you to specify the maximum usage rate for each HDD class in the 9900 subsystem. When a parity group exceeds this limit, the auto migration function makes a plan to migrate one or more volumes in this parity group to another parity group in a higher HDD class (see Figure 2.5). This storage tuning method addresses and can eliminate bottlenecks of disk drive activity and also provides load balancing of disk drive usage. HIHSM uses the estimated usage rates to verify each proposed auto migration, and will not perform a migration plan which might result in a target parity group exceeding the user-specified maximum disk usage rate.

Note: The default maximum usage values are rough estimates only and may not provide the best results. The user should specify the appropriate values for their operational environment.

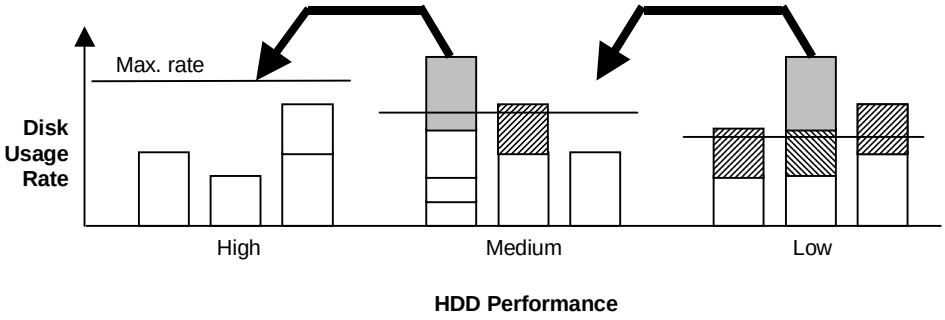

Figure 2.5 Migrating to Improve Disk Usage

Selecting source volumes. As shown in Figure 2.5, the HIHSM auto migration function identifies parity groups which exceed the user-specified usage limit, and then migrates the high-usage volumes to parity groups in higher HDD classes. When the parity groups in the highest HDD classes start to run out of reserved (empty) volumes, HIHSM automatically migrates low-usage volumes from higher HDD class groups to lower HDD class groups to maintain available reserve volumes (see Figure 2.6).

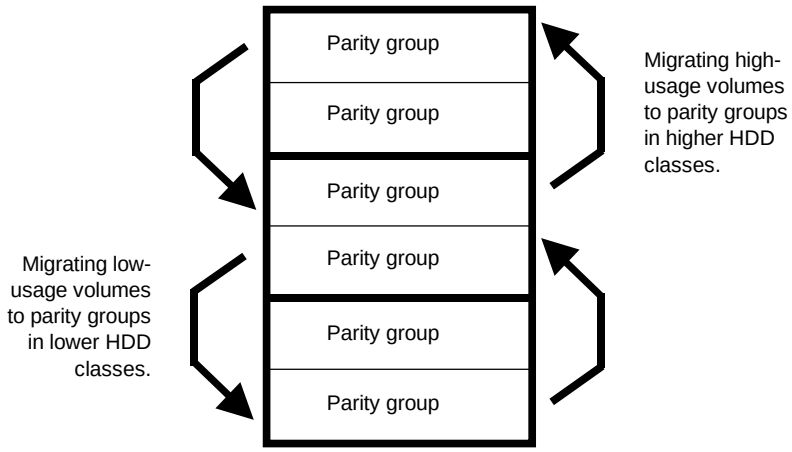

Figure 2.6 Keeping Reserve Volumes Available

Auto migration process. The HIHSM auto migration function performs the following process according to the user-specified schedule:

1. Analyze monitor data. You specify the monitor data term to be analyzed.
2. Make auto migration plan. You can view, modify, and cancel auto migration plans.
3. Perform auto migration plan. You specify when the auto migration plan is executed.
4. Analyze monitor data to confirm tuning results.

Figure 2.7 shows two examples of the auto migration process. In the first example, HIHSM uses the monitor data collected every day between 09:00 and 12:00 hours to make an auto migration plan, and executes the auto migration plan every day at 00:00 hours. In the second example, HIHSM uses the monitor data collected on every Tuesday between 17:30 and 22:00 hours to make an auto migration plan, and executes the plan at 23:00 hours. Figure 2.8 shows an example of the user modifying the auto migration plan.

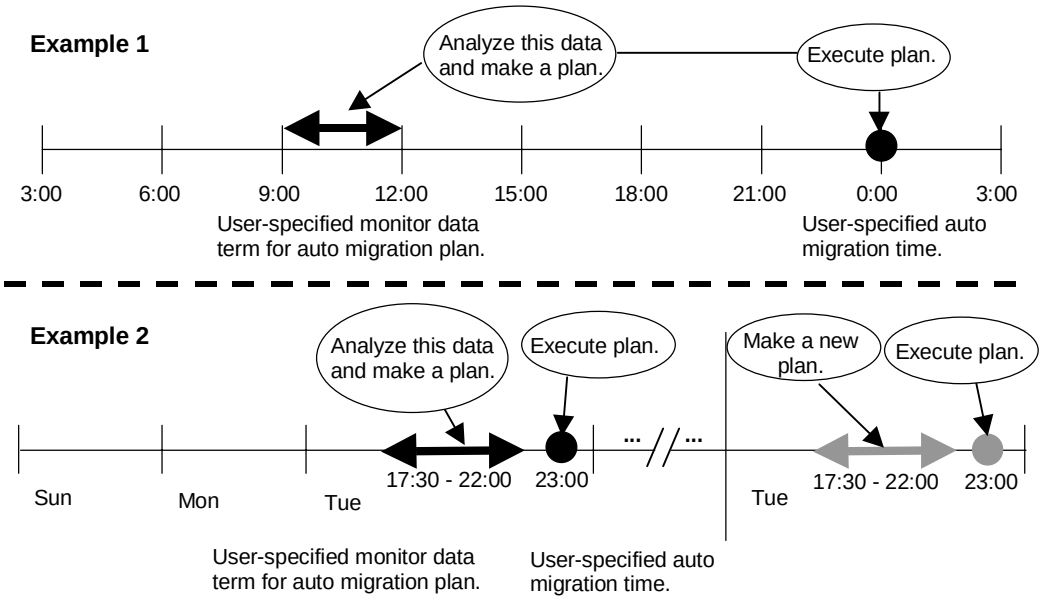

Figure 2.7 Making and Executing Auto Migration Plans

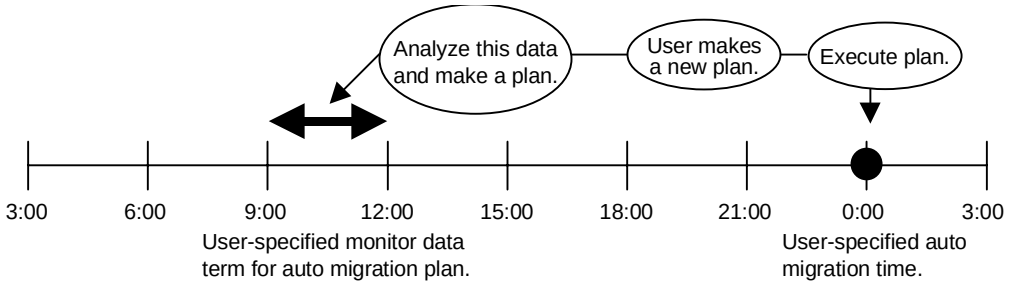

Figure 2.8 Modifying Auto Migration Plans

2.7 Manual Migration Operations

The HIHSM manual migration function enables you to take full control of Dynamic Optimizer tuning operations. Before performing manual migration operations, you can estimate the changes from the monitored information of your manual volume migrations and then determine the source and target volumes (i.e., make migration plans). Using HIHSM manual migration operations, you can fine-tune the 9900 subsystem to optimize parity group, volume, and back-end processor usage (DKPs, DRRs) for your operational environment.

While HIHSM auto migration operations are based on disk usage and hierarchy of parity groups, you can use manual migration operations to address back-end processor usage (DKPs and DRRs) as well as volume and parity group usage. If HIHSM monitoring shows high or unbalanced processor usage, you can use manual migration operations to tune the processor performance of the 9900 subsystem. Please refer to section 2.4 for further information on interpreting processor usage data and addressing overloaded/unbalanced processor resources.

2.8 GraphTool Software

The HIHSM GraphTool software enables you to extract the 9900 monitor data from the **artlog.lzh** export file (see section 4.4) and graph this statistical information in several data plotting formats. These graphical representations of 9900 monitor data enable you to identify workload characteristics and highlight key information about 9900 subsystem usage such as peaks and trends. Table 2.2 lists the types of data graphs provided by GraphTool. The GraphTool software is included as part of the HIHSM software product.

Table 2.2 Data Graphs

Name	Content
CHA Processor	CHP usage
DKA Processor	DKP usage
DRR	DRR usage
Access Paths	adapter-CSW, CSW-cache, and adapter-SM
Disk (Parity Group)	Parity group usage
Disk (Volume)	LDEV usage for the specified parity group

For instructions on installing and using the HIHSM GraphTool software, please see Chapter 5. GraphTool allows you to perform the following operations:

- Extract the data files from the compressed monitor data export files (**artlog.lzh**),
- Select the desired data file (9900 subsystem and month) to graph,
- Select the desired monitor data term (range of monitor data) to graph,
- Select the desired type of graph (i.e., plotting method), and
- View the specified monitor data on the specified graph.

Chapter 3 Installation Requirements and Procedures

3.1 System Requirements

Dynamic Optimizer operations involve the 9900 subsystem(s) and the licensed HIHSM remote console software. The system requirements for Dynamic Optimizer are:

- **9900 Remote Console PC.** The 9900 Remote Console PC can be attached to multiple 9900 subsystems via the 9900-internal LAN. Please refer to the *9900 Remote Console User's Guide* for instructions on installing and using the 9900 Remote Console PC and Remote Console Main (RMCMAIN) software.

  Note: Administrator access to the RMCMAIN software is required to perform HIHSM operations. Users without administrator access can only view HIHSM information.

- **RMCMAIN and DKCMAIN license keys for HIHSM.** The HIHSM remote console software is installed using the HIHSM RMCMAIN license key. You must have a separate HIHSM DKCMAIN license key for each 9900 subsystem on which you want to perform HIHSM operations.

3.2 Installing the Hardware

Installation of the 9900 hardware is performed by the user and the Hitachi Data Systems representative. To install the hardware required for HIHSM operations:

1. **User**: Identify the 9900 subsystems on which you want to perform HIHSM operations, and install the 9900 Remote Console PC near these subsystems. Please refer to the *9900 Remote Console User's Guide* for instructions on installing the Remote Console PC.
2. **Hitachi Data Systems Representative**: Connect the 9900 Remote Console PC to the 9900 subsystems via the 9900-internal LAN. Please refer to the *9900 Remote Console User's Guide* for further information on Remote Console installation and configuration.

3.3 Enabling the HIHSM Feature

The user enables the HIHSM option on the Remote Console PC and the HIHSM option on each 9900 subsystem using the RMCMAIN and DKCMAIN license key codes for HIHSM. *Note:* The RMCMAIN and DKCMAIN license key codes for a 9900 subsystem are identical. However, you must have separate DKCMAIN license key codes for each 9900 subsystem. You may not re-use the same DKCMAIN key code for multiple 9900 subsystems.

To enable the HIHSM feature:

1. Check with your Hitachi Data Systems representative to verify that the correct microcode and SVP software are installed and enabled on the 9900 subsystems which will perform HIHSM operations. Also make sure that your RMCMAIN software version is correct.
2. Make sure that the 9900 Remote Console PC and RMCMAIN software are installed and functioning properly. Refer to the *9900 Remote Console User's Guide* for instructions on installing the Remote Console PC and RMCMAIN software.
3. Enable the remote HIHSM option on the Remote Console PC as follows:
   a) Start up and log in to the 9900 RMCMAIN software with administrator access.
   b) Select **Option…** to open the RMCMAIN Option Product panel (see Figure 3.1).
   c) On the Option Product panel, select **Remote HIHSM**, and then select **Install…** to open the Input Key Code panel (see Figure 3.2).
   d) Enter the license key code in the **Key Code** text box, and then select **OK**.
   e) If the key code is accepted, the Program Product (P.P.) Confirmation panel opens (see Figure 3.3). Confirm the information displayed on this panel, and select **Install**. The Option Product panel now displays **[Install]** for the **Remote HIHSM** option.
   f) Select **Close** on the Option Product panel to return to the Remote Console Main panel.
4. If not already done, add the attached 9900 subsystems to the Remote Console PC. Select **Controller…**, select **Add…**, enter the subsystem name and serial number, and select **OK**. Then select the subsystem you just added on the Connection Control panel, and select **Entry**. Refer to the *9900 Remote Console User's Guide* for more detailed instructions.
5. Enable the HIHSM option on each subsystem as follows:
   a) On the Connection Control panel, select the desired 9900 subsystem, and then select **Install…** to open the DKCMAIN Option Product panel (see Figure 3.4).
   b) Select **HIHSM**, select **Install…**, enter the license key code for the selected subsystem on the Input Key Code panel, and select **OK**.
   c) Confirm the information displayed on the P.P. Confirmation panel, and select **Install** to enable the selected HIHSM option on the selected subsystem. The DKCMAIN Option Product panel now displays **[Install]** for the selected HIHSM option.
6. After enabling the HIHSM option on the Remote Console PC and on the 9900 subsystems, you are now ready to start performing HIHSM operations.

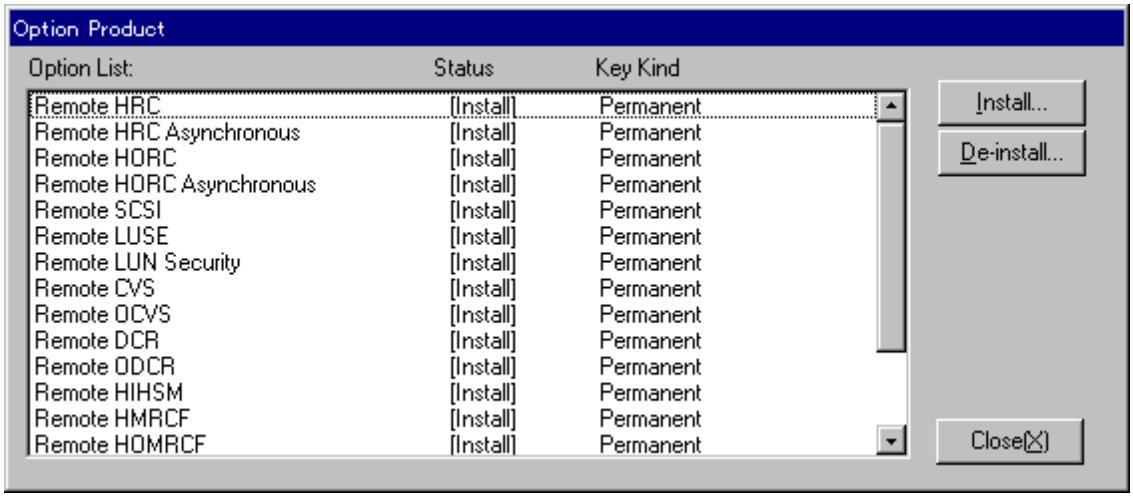

Figure 3.1 Enabling the Remote HIHSM Option

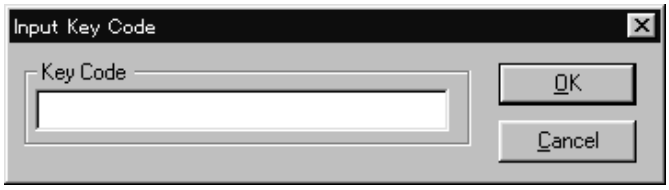

Figure 3.2 Entering the HIHSM License Key Code

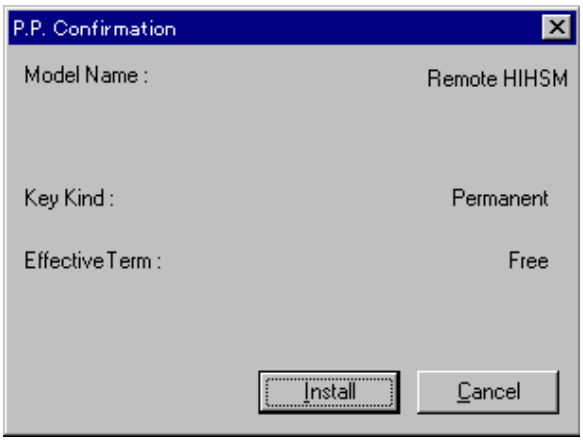

Figure 3.3 Confirming the HIHSM Key Code

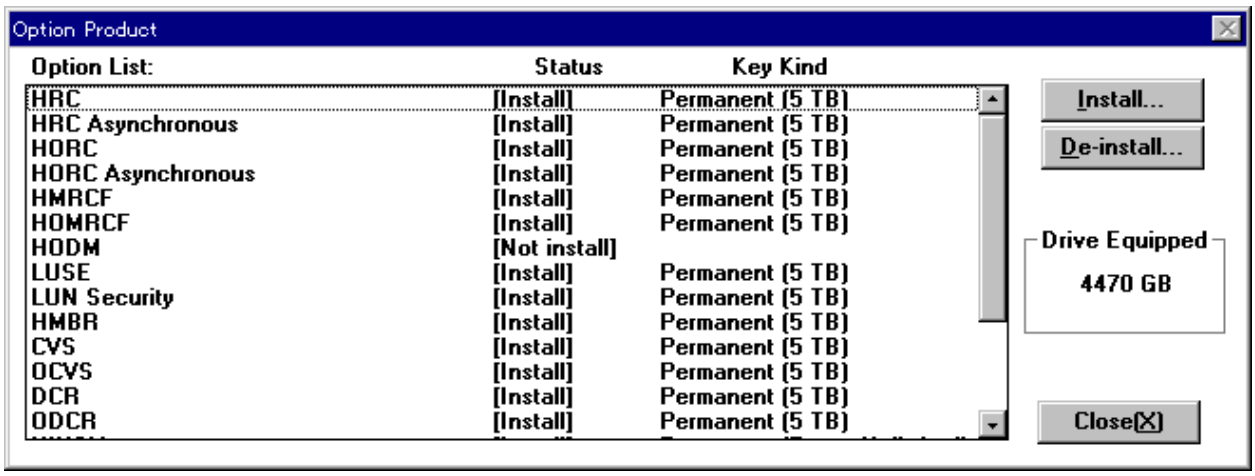

Figure 3.4 Enabling the HIHSM Option on Each Subsystem

Chapter 4 Performing Dynamic Optimizer Operations

HIHSM migration operations should only be performed when you can expect a large improvement in 9900 subsystem performance. HIHSM migration may not provide significant improvement for cases in which parity group or volume usage varies only slightly, or for cases in which overall DKP or DRR usage is relatively high. Also keep in mind that HIHSM migration operations can improve subsystem performance in one area while at the same time decreasing performance in another area. For example, let's say that parity groups A and B have average usage values of 20% and 90%, respectively. HIHSM estimates that if one logical volume is migrated from parity group B to parity group A, the usage values will become 55% and 55%. If this migration operation is performed, the I/O response time for parity group B will probably decrease, but the I/O response time for parity group A may increase, which may or may not improve overall throughput.

The HIHSM remote console software allows you to control HIHSM monitoring, reserve volumes as HIHSM target volumes, configure HIHSM auto migration operations, export the 9900 monitor data, and view the history logs of HIHSM operations. The GraphTool software (see Chapter 5) is used to extract the monitor data from the export file and graph the data.

HIHSM operations include:

- Starting the HIHSM remote console software (see section 4.1),
- Using the HIHSM Main panel (see section 4.2),
- Monitoring 9900 subsystem usage (see section 4.3),
- Exporting monitor data (see section 4.4),
- Viewing monitor data (see section 4.5),
- Reserving target volumes (see section 4.6),
- Configuring auto migration operations (see section 4.7),
- Starting, checking, and stopping migration operations (see section 4.8), and
- Viewing the main history log (see section 4.9).

Please see Chapter 5 for instructions on installing the GraphTool software and using GraphTool to extract and graph the data in a HIHSM export file.

4.1 Starting the HIHSM Remote Console Software

As soon as HIHSM hardware and software installation are complete (see Chapter 3), you can start the HIHSM remote console software. To start the HIHSM software:

1. Start up and log in to the 9900 RMCMAIN software with administrator access or custom HIHSM access. If you do not have administrator access or custom HIHSM access, you will only be able to view the HIHSM information for the connected subsystem.
2. Connect to the desired 9900 subsystem (refer to the *9900 Remote Console User's Guide* for instructions). Remember that the remote HIHSM feature must be enabled on each 9900 subsystem on which you want to perform HIHSM operations (see section 3.3).
3. When you have connected to the selected subsystem, the Option Select panel is displayed. Select the **HIHSM** button to start the HIHSM remote console software.
4. HIHSM now asks if you want to gather (load) the statistical information (monitor data) from the subsystem (DKC = disk controller) (see Figure 4.1). If you have already started HIHSM monitoring, select **OK** to load the monitor data. If you have not yet started HIHSM monitoring, select **Cancel** to go directly to the HIHSM Main panel.
5. If you selected **OK**, HIHSM loads the monitor data (see Figure 4.2) and then opens the HIHSM Main panel. If you selected **Cancel**, the HIHSM Main panel opens right away.

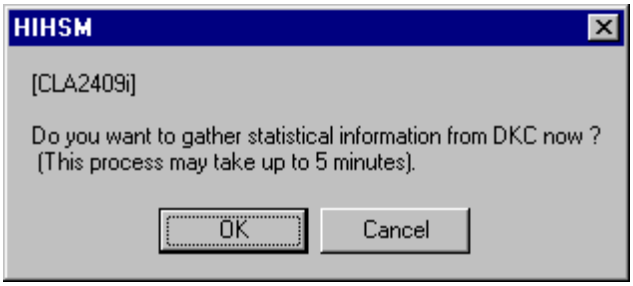

Figure 4.1 Launching HIHSM

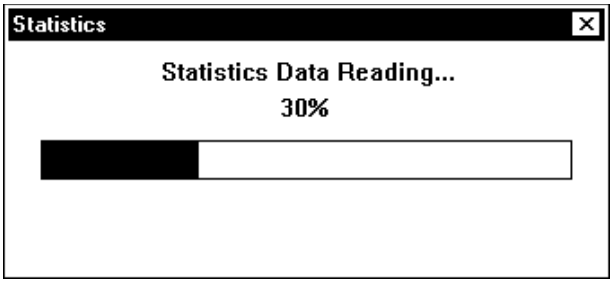

Figure 4.2 Loading the Monitor Data

4.2 Using the HIHSM Main Panel

The HIHSM Main panel (see Figure 4.3) displays the monitor data term and parity group information for the connected 9900 subsystem and provides access to all HIHSM operations. To open the HIHSM Main panel, select **HIHSM** on the RMCMAIN Option Select panel.

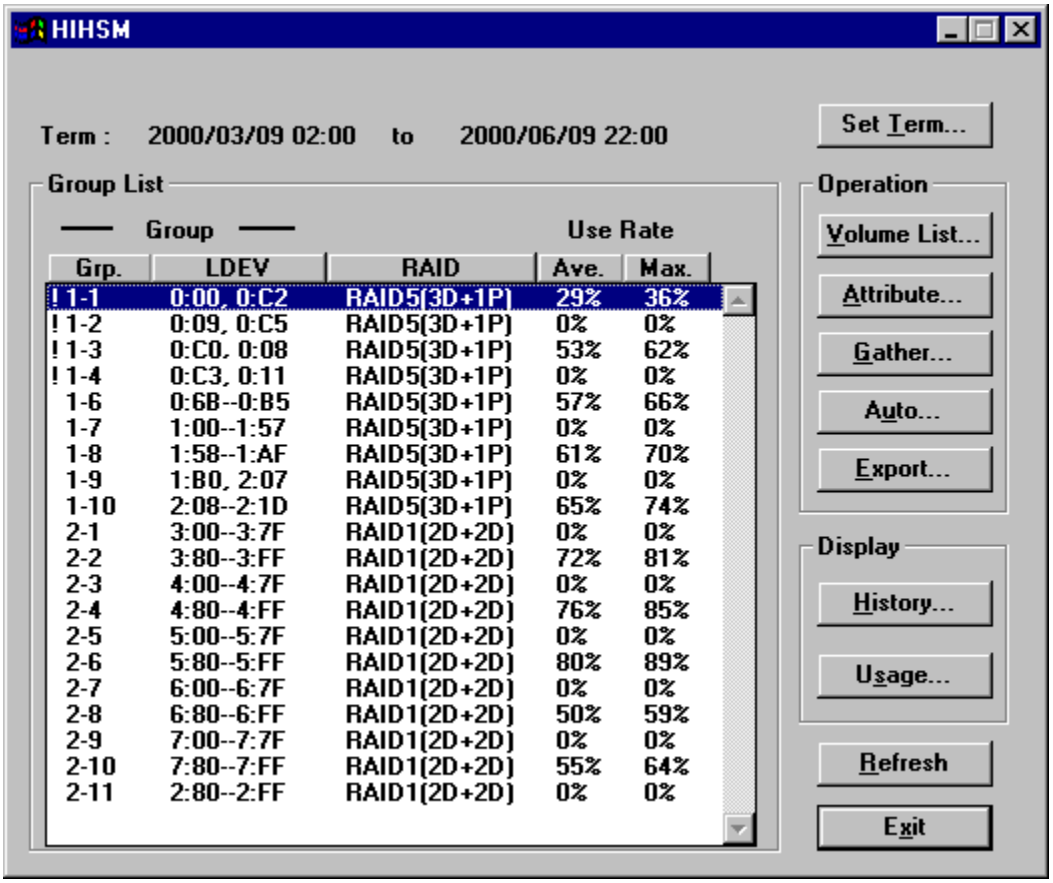

Figure 4.3 HIHSM Main Panel

The **Term** field displays the monitor data term currently being displayed. Each time you open the HIHSM main panel, the information of the monitor data term set in the previous session is displayed until you refresh the panel. When you open the panel for the first time, all available information (maximum = three months) is displayed. The Data Term panel (see section 4.3.2) allows you to change the data term being displayed. To open the Data Term panel, select the **Set Term...** button.

The **Group List** box displays the following information for the connected 9900 subsystem:

- **Grp**. Installed parity groups by group ID: [frame number] - [parity group number]. If the precision of the parity group usage rate is likely to be low (e.g., volumes in the group have been migrated by HIHSM, or the configuration has changed by CVS), an exclamation mark (!) is displayed before the parity group ID (at the far left end of the row).

- **LDEV**. Logical CU number and device ID (two-digit hexadecimal value 00-FF).
- **RAID**. RAID configuration: RAID5 or RAID1.
- **Use Rate - Ave***. The average usage rate for the parity group during the data term.
- **Use Rate - Max***. The maximum usage rate for the parity group during the data term.

* After the volume migration, the precision of the old data (before migration) will become low. The usage rate is indicated as follows:
"+0%": The maximum usage rate is not zero, but the average usage rate is zero.
"00%": Both the maximum and average usage rates are zero.
"--%": The usage rate cannot be calculated.

The **Set Term…** button opens the Data Term panel (see section 4.3.2), which allows you to set the monitor data term to display on the HIHSM Main panel.

The **Operation** box provides access to the following HIHSM operations:

- The **Volume List…** button opens the Volume List panel (see section 4.8), which displays the detailed volume information for the selected parity group and allows you to start a manual migration operation, stop a manual or automatic migration operation, and check the status of any in-process migration operation.
- The **Attribute…** button opens the Change Attribute panel (see section 4.6), which allows you to reserve (or unreserve) volumes as HIHSM target volumes.
- The **Gather…** button opens the Data Gathering panel (see section 0), which allows you to set or reset subsystem monitoring.
- The **Auto…** button opens the Auto Migration panel (see section 4.7), which allows you to configure and perform automatic volume migration.
- The **Export…** button opens the Export Monitor panel (see section 4.4), which allows you to export monitor data to a compressed file on diskette.

The **Display** box provides access to the following HIHSM operations:

- The **History…** button opens the History panel (see section 4.9), which displays the HIHSM main history log for all migration operations. The auto migration history log is accessed from the Auto Migration panel.
- The **Usage…** button opens the Usage panel (see section 4.5.3), which displays the detailed processor usage rates (CHPs, DKPs, and DRRs), access path usage rates, and write-pending rates.

The **Refresh** button updates the information displayed on the Group List box.

Caution: You should always confirm the latest status (configuration information) on the HIHSM Main panel before and after migration.

4.3 Monitoring 9900 Subsystem Usage

The HIHSM remote console software allows you to monitor 9900 subsystem performance and collect detailed usage statistics for the 9900 processors, access paths, parity groups, and logical volumes. HIHSM monitor operations include:

- Starting and stopping subsystem monitoring (see section 4.3.1),
- Selecting the monitor data term (see section 4.3.2), and
- Saving monitor data to disk (see section 4.4).

Caution: When an error condition exists in the 9900 subsystem, resource usage can increase or become unbalanced. Do not use 9900 monitor data collected during an error condition as the basis for planning HIHSM migration operations.

4.3.1 Starting and Stopping Monitoring

The Data Gathering panel (see Figure 4.4) allows you to start/stop HIHSM monitoring and specify the start time for HIHSM monitoring. The HIHSM monitoring function collects 9900 usage data for the 23-hour period starting at the user-specified start time. To open the Data Gathering panel, select the **Gather…** button on the HIHSM Main panel.

Caution: Make sure to schedule your HIHSM monitoring around the 9900 SVP auto-reboot time. When the SVP reboots, HIHSM monitoring is discontinued and will not resume until the next scheduled start time.

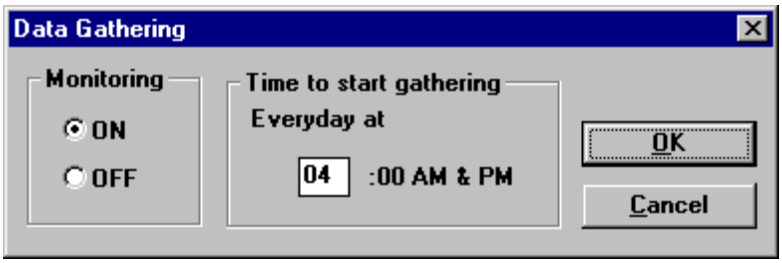

Figure 4.4 Data Gathering Panel

The **Monitoring** box allows you to turn on (**ON**) and turn off (**OFF**) the HIHSM monitoring function. The **Time to start gathering** box allows you to select the time (00:00-11:00 except 01:00, 02:00 and 03:00) to start collecting monitor data. For example, if you turn monitoring **ON** and enter **04:00**, the SVP will collect and store HIHSM monitor data twice per day starting at 4 am (04:00 hours) and 4 pm (16:00 hours).

To start collecting HIHSM monitor data for the connected 9900 subsystem:

1. On the HIHSM Main panel, select the **Gather…** button to open the Data Gathering panel.
2. On the Data Gathering panel, select the **ON** button, and then enter the desired start time for HIHSM monitoring.
3. Select **OK** to start collecting 9900 monitor data each day at the specified time.

Caution: When you turn HIHSM monitoring OFF, the monitor data which has not yet been loaded by the HIHSM software will be discarded. If you need this data, exit and restart HIHSM to load the remaining monitor data, and then turn off the monitoring function.

To stop collecting HIHSM monitor data for the connected 9900 subsystem:

1. On the HIHSM Main panel, select the **Gather…** button to open the Data Gathering panel.
2. Select **OFF** to turn HIHSM monitoring off, and then select **OK**.

4.3.2 Selecting the Monitor Data Term

The Data Term panel (see Figure 4.5) allows you to select the range of the latest 9900 monitor data collected by SVP (monitor data term) to be displayed on the HIHSM Main panel and Usage panel. The processor, access path, parity group, and volume usage rate values are calculated using the data in the specified monitor data term. For manual migration, HIHSM uses the data in the specified monitor data term to calculate estimated usage rates. To open the Data Term panel, select **Set Term…** on the HIHSM Main panel.

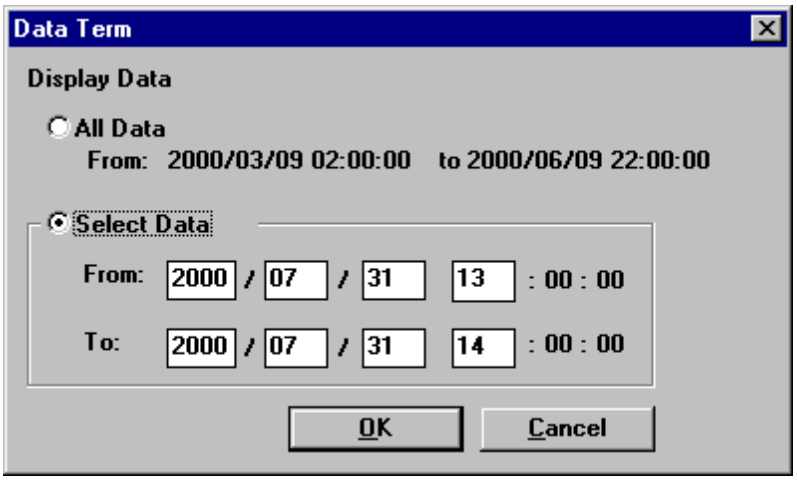

Figure 4.5 Selecting the Monitor Data Term

To select the monitor data term to be displayed:

1. On the HIHSM Main panel, select the **Set Term…** button to open the Data Term panel.
2. On the Data Term Panel, select **All Data** to display all monitor data, or select **Select Data** and enter the starting date/time (**From**) and ending date/time (**To**) for the desired range of monitor data.
3. Select **OK** to display the selected monitor data term (all or specified range).

Note: If you specify a large monitor data term (for **Select Data**), the display speed of the screens related to manual migration operations will decrease. Therefore, you should specify the range as narrow as possible to maintain a normal screen display speed.

4.4 Exporting Monitor Data

The Export Monitor panel (see Figure 4.6) allows you to export 9900 monitor data to a compressed file on floppy diskette (**a:\artlog.lzh**). The HIHSM GraphTool software (see Chapter 5) allows you to extract the exported monitor data from this file and graph the data. To open the Export Monitor panel, select the **Export…** button on the HIHSM Main panel.

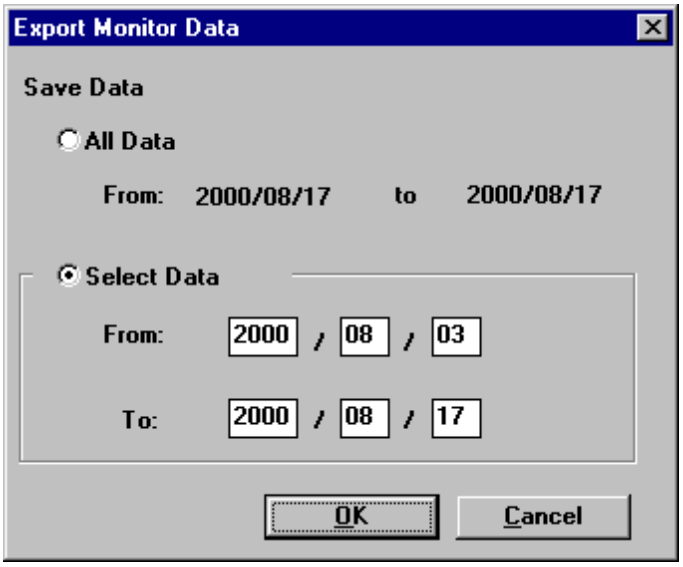

Figure 4.6 Export Monitor Panel

To export 9900 monitor data to disk:

1. Put a (blank) floppy diskette in drive **a:** of the Remote Console PC.
2. On the HIHSM Main panel, select the **Export…** button to open the Export Monitor panel.
3. On the Export Monitor panel, select **All Data** to save all monitor data to disk, or select **Select Data** and then enter the desired range of monitor data to save to disk.
4. Select **OK** to export the selected monitor data to the floppy disk drive. HIHSM creates a compressed file called **artlog.lzh** on the **a:** drive.
5. See Chapter 5 for instructions on extracting the monitor data from this export file and graphing the monitor data.

4.5 Viewing Monitor Data

HIHSM displays the parity group, logical volume, processor, access path, and write-pending rate information collected by the HIHSM monitor function. The usage rate values are based on the monitor data collected during the user-specified monitor data term (see section 4.3.2).

4.5.1 Viewing Parity Group Usage Rates

The HIHSM Main panel (see Figure 4.7) displays the average and maximum usage rates for each parity group in the 9900 subsystem. The parity group usage value is the sum of the logical volume usage values for all volumes in the parity group.

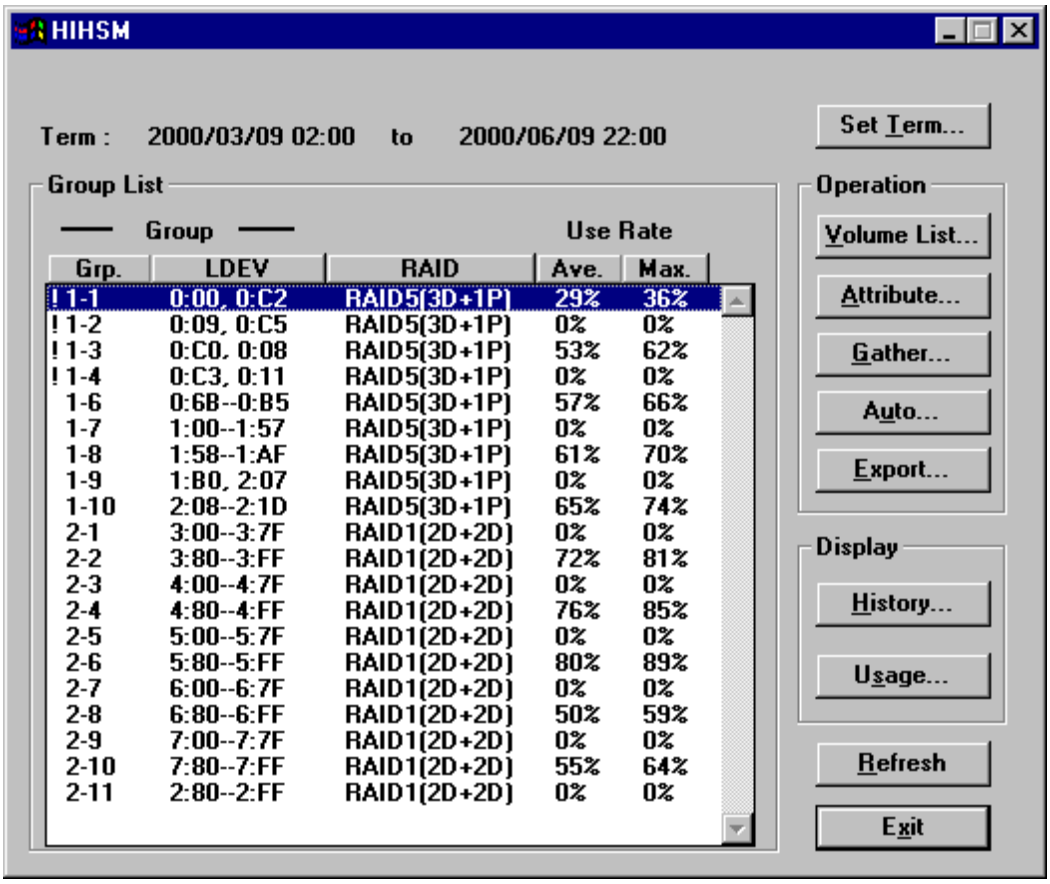

Figure 4.7 Viewing Parity Group Usage

If the HIHSM monitor data shows overall high parity group usage, you should consider installing additional HDDs, and then migrating the high-usage volumes to the new parity groups. If HIHSM monitoring shows unbalanced parity group usage, you should use HIHSM auto migration to balance parity group usage.

4.5.2 Viewing Volume Usage Rates

The Volume List panel (see Figure 4.8) displays the average and maximum usage rates (including sequential and random access) for each logical volume in the selected parity group. To open the Volume List panel, select the desired parity group on the HIHSM Main panel, and then select the **Volume List…** button.

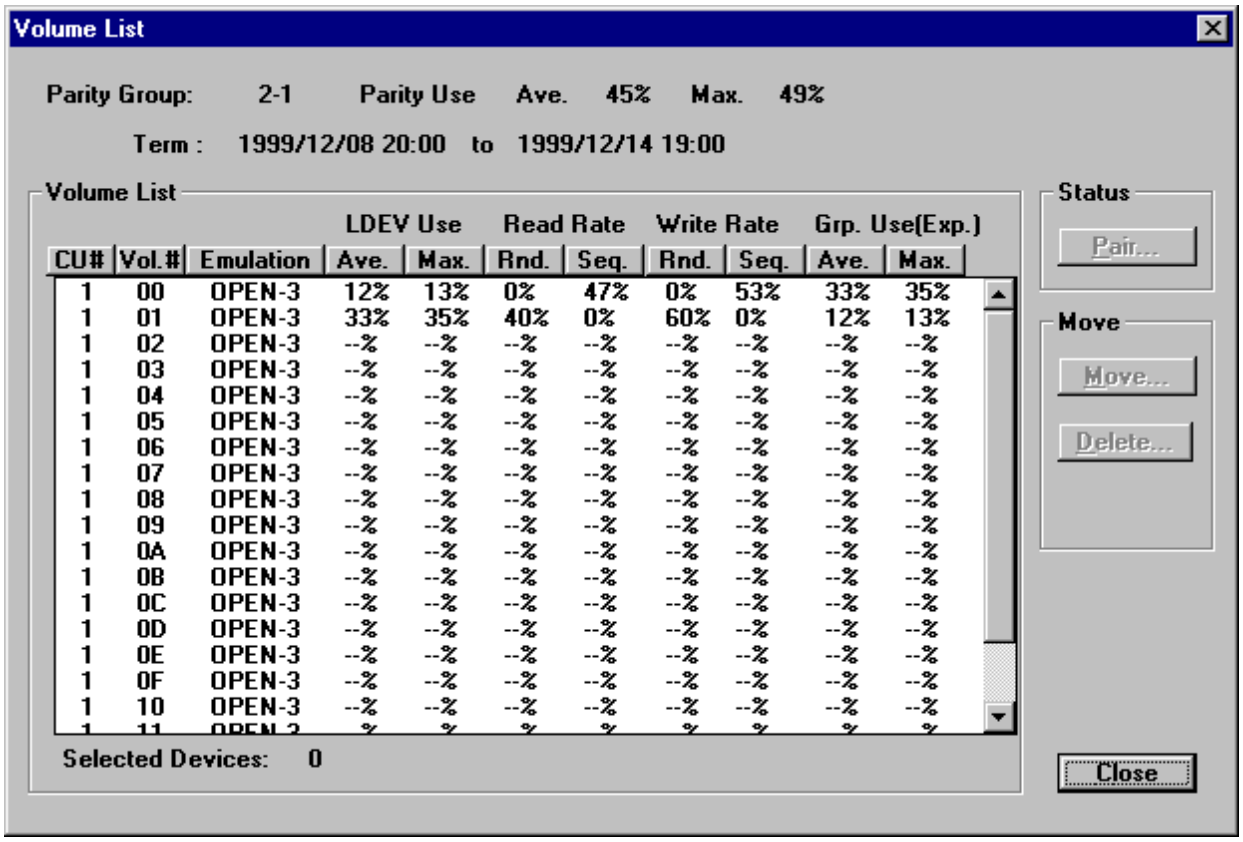

Figure 4.8 Viewing Volume Usage

The **Parity Group** field at the top of the Volume List panel displays the ID (frame number - group number) and usage rates of the selected parity group. The **Term** field displays the current monitor data term.

The **Volume List** box displays the following information for each logical volume in the selected parity group:

- **CU#**: Logical CU number (0, 1, 2, … F). If the precision of the volume usage rate is likely to be low, an exclamation mark (!) is displayed before the CU number (at the far left end of the row). If a volume is being migrated, an asterisk (*) is displayed before the CU number.
- **Vol#**: LDEV ID (00-FF).
- **Emulation**: Device emulation type (LVI or LU type).
- **LDEV Use**: Average usage rate and maximum usage rate.
- **Read Rate**: Random (**Rnd**) read and sequential (**Seq**) read I/O usage rates.
- **Write Rate**: Random (**Rnd**) write and sequential (**Seq**) write I/O usage rates.
- **Grp.Use(Exp.)**: Expected (estimated) average and maximum usage rates of the parity group, if the volume were migrated out of the group (or deinstalled).

4.5.3 Viewing Processor and Access Path Usage Rates

The Usage panel (see Figure 4.9) displays the detailed processor and access path usage statistics for the specified monitor data term. To open the Usage panel, select the **Usage…** button on the HIHSM Main panel.

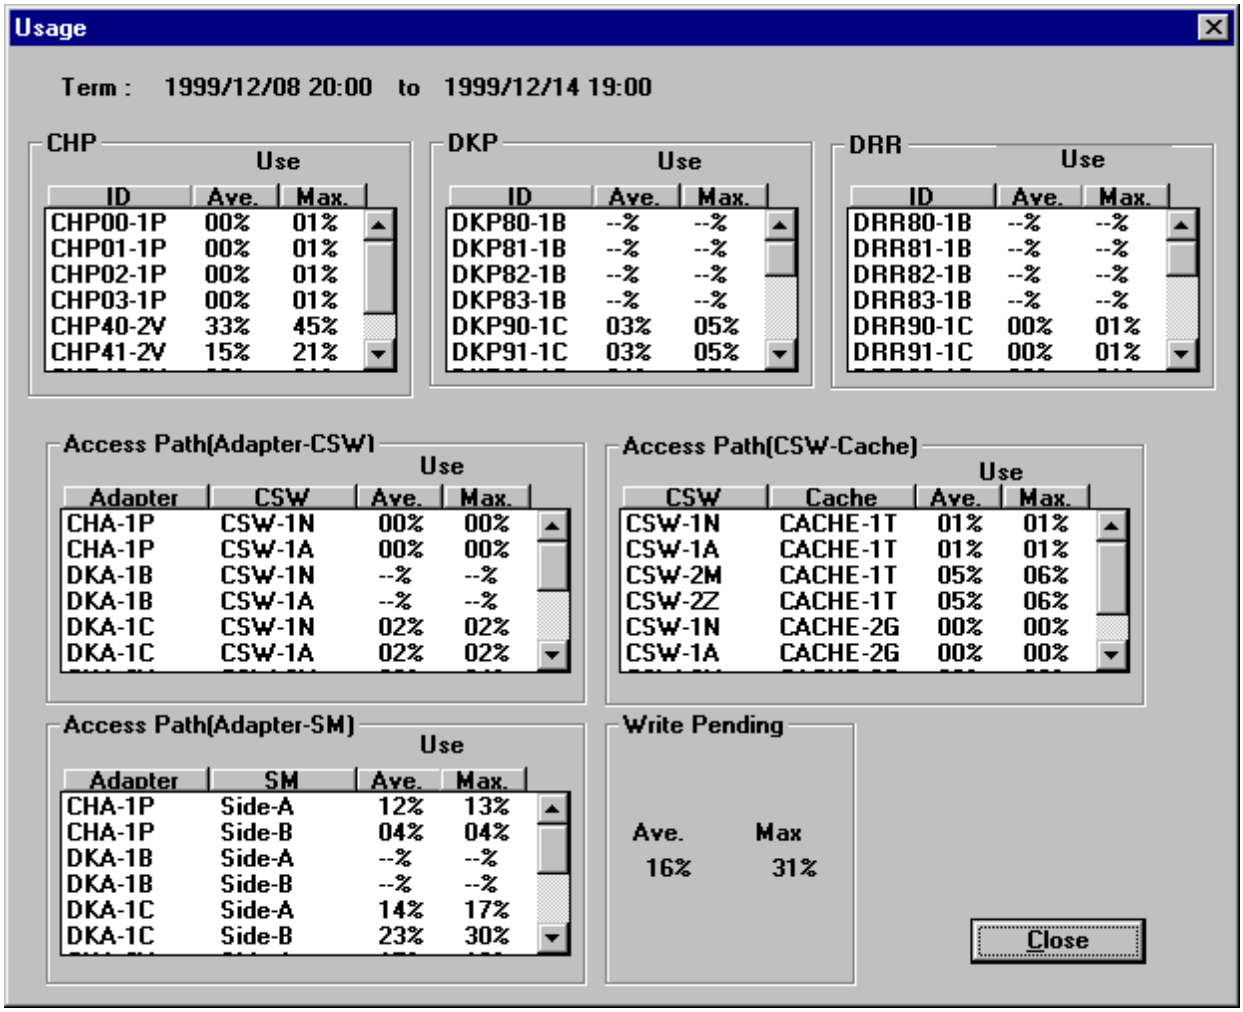

Figure 4.9 Viewing Processor and Access Path Usage

The Usage panel displays the following information for the selected parity group:

- **CHP:** Average and maximum usage of the CHPs for the specified monitor data term.
- **DKP:** Average and maximum usage of the DKPs for the specified monitor data term.
- **DRR:** Average and maximum usage of the DRRs (data recovery and regeneration processors) for the specified monitor data term.
- **Access Path:** Average and maximum usage of the control and data access paths (adapter-CSW, CSW-cache, and adapter-SM) for the specified monitor data term.
- **Write Pending:** Average and maximum write-pending rate for the specified data term.

Note: HIHSM does not estimate processor or access path usage rates. Please refer to section 2.4 for information on interpreting processor and access path usage rates and addressing overloaded or unbalanced processor resources.

4.6 Reserving Target Volumes

The **Normal** and **Reserved** volume attributes indicate which volumes are available for use as HIHSM target volumes. HIHSM target volumes must be reserved. A normal volume cannot be used as a target volume. Before HIHSM migration is performed, you need to decide which volumes will be the HIHSM target volumes and change their volume attribute to **Reserved**.

The Change Attribute panel (see Figure 4.10) displays and allows you to change the reserve attribute of the LDEVs in the selected parity group. To open the Change Attribute panel, select a parity group on the HIHSM Main panel, and then select the **Attribute…** button.

To reserve or unreserve a volume:

1. Select the desired parity group on the HIHSM Main panel, and then select the **Attribute…** button to open the Change Attribute panel.
2. On the Change Attribute panel, select the desired volume(s), and then select **Normal** or **Reserved**. Only one of these buttons is enabled, based on the selected volume(s).
   Note: If a volume is not displayed on this panel, it is already being used for another function (e.g., command device for CCI operations).

   Note: The volumes are displayed by logical CU image and LDEV ID (not SCSI path).
3. The Change Attribute Confirmation panel (see Figure 4.11) displays your requested change. To delete a volume from the list, select it and then select **Omit**. To cancel your request, select **Cancel**. To make the requested attribute change(s), select **OK**.

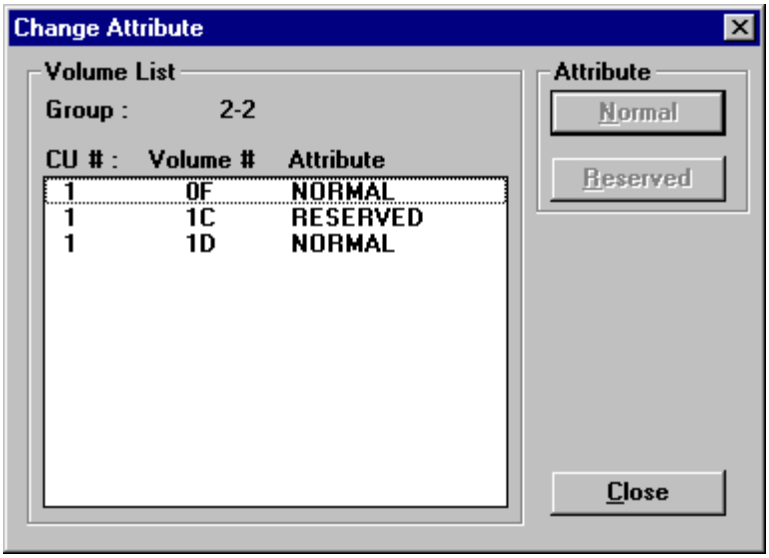

Figure 4.10 Change Attribute Panel

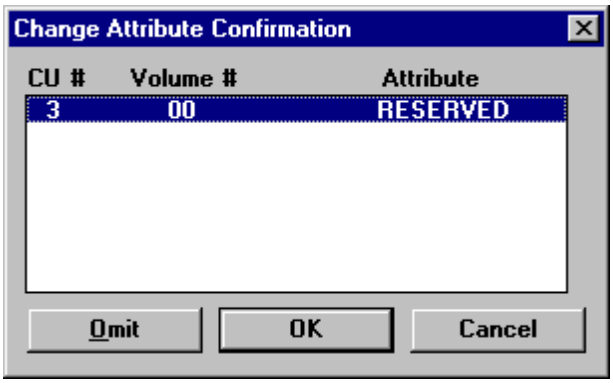

Figure 4.11 Reserving/Unreserving a Volume

4.7 Configuring Auto Migration Operations

The HIHSM auto migration function takes the range of 9900 monitor data specified by the user, analyzes parity group utilization against the user-specified maximum disk usage rates, and creates and executes auto migration plans to improve the disk drive and parity group utilization of the 9900 subsystem. The auto migration function uses the estimated usage rates to verify its auto migration plans. If any estimated usage rate is over the user-specified maximum rate, HIHSM will not make the auto migration plan.

The Auto Migration panel (see Figure 4.12) displays the auto migration plan information and provides access to all auto migration functions. To open the Auto Migration panel, select the **Auto…** button on the HIHSM Main panel.

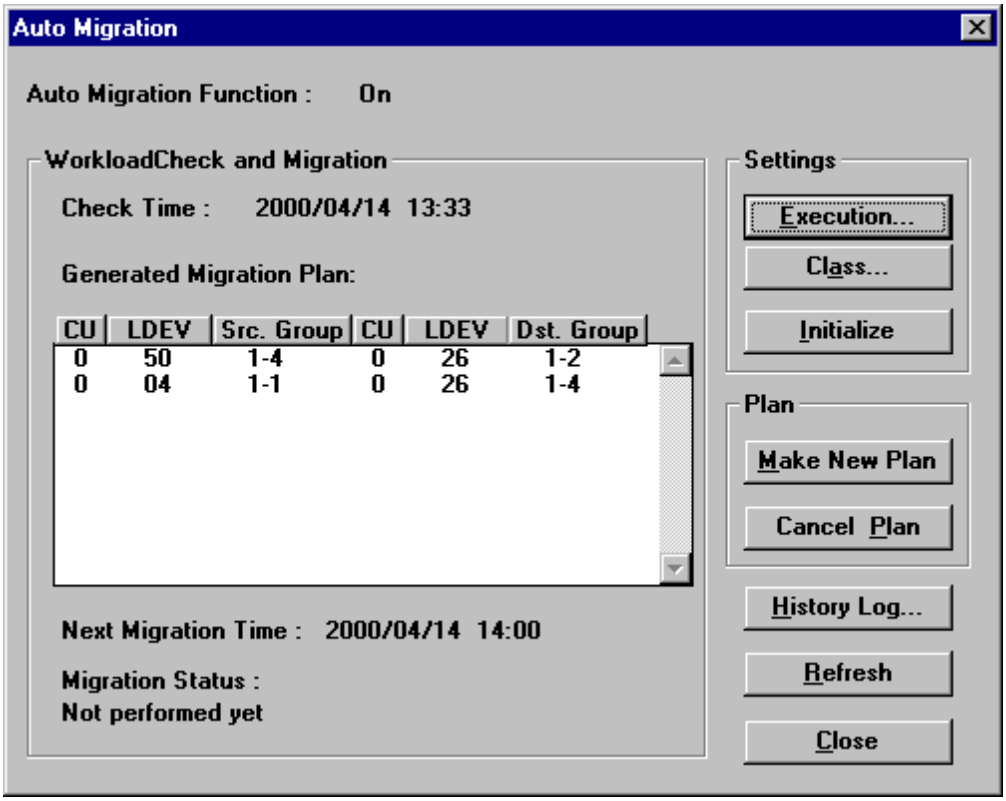

Figure 4.12 Auto Migration Panel

The Auto Migration panel displays the following information:

- **Auto Migration Function:** Status (**On** or **Off**) of the auto migration function. To turn on/off the auto migration function, select the **Execution…** button to open the Execution Parameter Settings panel (see section 4.7.3).
- **Workload Check and Migration:**
  - **Check Time:** The next time* that the auto migration function will make a plan.

– **Generated Migration Plan:**

CU	CU image number of the source volume.
LDEV	LDEV ID of the source volume.
Src. Group	Parity group ID of the source volume.
CU	CU image number of the target volume.
LDEV	LDEV ID of the target volume.
Dst. Group	Parity group ID of the target volume.

– **Next Migration Time:** The next time* that the auto migration function will execute an auto migration plan. If HIHSM has not made a plan yet, "--" is displayed.

– **Migration Status:**

Not planned yet.	HIHSM has not yet made a plan.
Not performed yet.	HIHSM has made a plan but has not yet executed it.
Failed to make plan.	HIHSM could not make the plan.
Under migration.	HIHSM is currently executing the plan.
Last migration has canceled. (Please see log file).	HIHSM canceled and discarded the migration plan.
Migration successfully ended. Plan has done.	HIHSM successfully executed the plan.

**Note:* The time values displayed on this panel are retrieved from the 9900 SVP, and may not correspond to the time displayed on the 9900 Remote Console PC.

The **Settings** box contains the following buttons:

- The **Execution…** button opens the Execution Parameter Settings panel (see section 4.7.3), which allows you to set the detailed parameters for the auto migration function.
- The **Class…** button opens the Class List panel (see section 4.7.1), which displays the HDD class assignment for each installed parity group and provides access to the fix parity group and set maximum disk usage functions. For further information on HDD class types, please refer to section 2.1.
- The **Initialize** button resets all auto migration settings/parameters to their default values. *Note:* Select this button after changing the configuration.

The **Plan** box contains the following buttons:

- The **Make New Plan** button allows you to delete the existing auto migration plan and make a new (different) auto migration plan. *Caution:* Do not select this button when auto migration operations are in progress.
- The **Cancel Plan** button cancels and discards the existing auto migration plan. HIHSM will make a new plan at the next scheduled check time. *Note:* You cannot cancel a migration plan which is currently being executed. If you need to cancel an in-process auto migration operation, you can delete the HIHSM pair (see section 4.8.3).

The **History Log…** button opens the Automatic Migration History Display panel (see section 4.7.6), which displays the auto migration history log. The **Refresh** button updates the auto migration status information displayed on the Auto Migration panel. The **Close** button closes the Auto Migration panel and returns you to the HIHSM Main panel.

4.7.1 Setting Fixed Parity Groups

A fixed parity group is a parity group that cannot be used by the auto migration function. If you want to exclude all volumes in a parity group from auto migration operations, you need to set the parity group as fixed before configuring and starting auto migration operations. Volumes in fixed parity groups will not be used for auto migration operations, but can be used as needed for manual migration operations.

The Class List panel (see Figure 4.13) lists the classes of installed HDDs and displays the number of parity groups, number of volumes, and number of reserved volumes in each class. To open the Class List panel, select the **Class…** button on the Auto Migration panel.

Note: Users should use the parameter settings with default values to apply the latest configuration information data. (Use the **Initialize** button in the Auto Migration panel.)

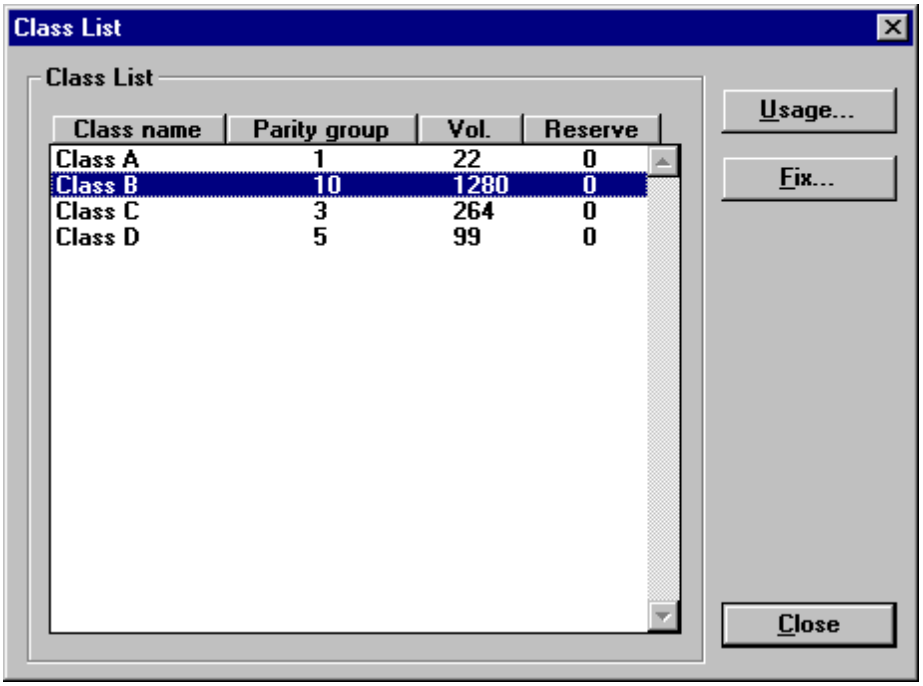

Figure 4.13 Class List Panel

To set or reset a fixed parity group:

1. On the HIHSM Main panel, select the **Auto…** button to open the Auto Migration panel.
2. On the Auto Migration panel, select the **Class…** button to open the Class List panel.
3. After reviewing the information displayed on the Class List panel, select the **Fix…** button to open the Fix Parity Group panel (see Figure 4.14).

4. The Fix Parity Group panel displays a list of fixed parity groups and a list of normal (not fixed) parity groups and allows you to add groups to and delete groups from each list.
   - To add a parity group to the list of fixed parity groups, select the parity group in the **Parity Group** list, and then select the **←Add** button.
   - To remove a parity group from the list of fixed parity groups, select the parity group in the **Fixed Parity Group** list, and then select the **Del→** button.
5. When the **Fixed Parity Group** list contains the desired parity groups, select **OK** to set/reset the parity groups as specified. You are returned to the Class List panel.

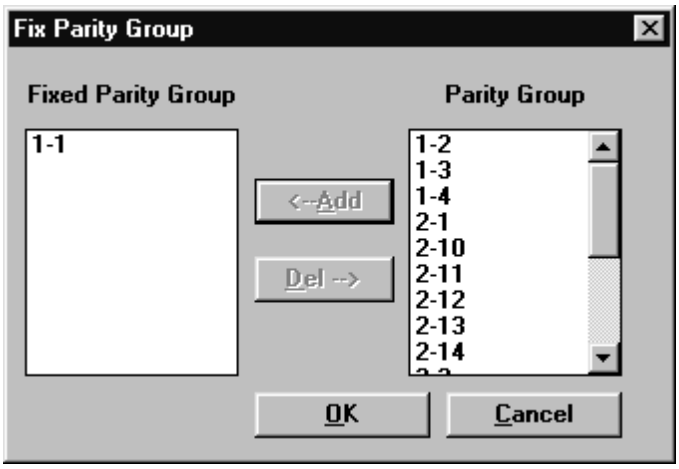

Figure 4.14 Fix Parity Group Panel

4.7.2 Setting the Maximum Disk Usage

The Class Usage panel (see Figure 4.15) displays the parity groups in the selected HDD class and allows you to set the maximum disk usage rate for the class. To open the Class Usage panel, select the desired class on the Class List panel, and select the **Usage…** button.

Note: The default maximum usage values are rough estimates only and may not provide the best results. The user should specify the appropriate values for their operational environment.

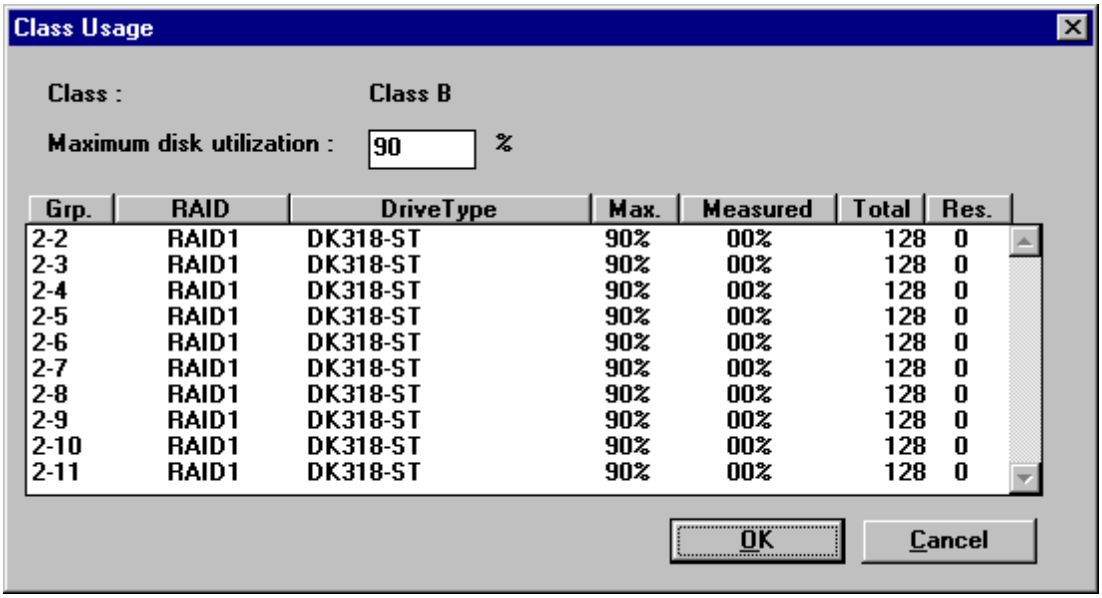

Figure 4.15 Class Usage Panel

The Class Usage panel displays the following information:

- **Class:** class type of the selected HDD class.
- **Maximum disk utilization:** allows you to specify the maximum disk usage for this class.
- Parity group information for the selected HDD class:
  - **Grp.:** Parity group ID (frame number - group number).
  - **RAID:** RAID configuration (RAID1 or RAID5).
  - **Drive Type:** HDD type. Table 4.1 lists the HDD types and the capacity of each type.
  - **Max.:** Maximum usage value (as specified on this panel).
  - **Measured:** Measured usage value for the specified monitor data term.
    Note: "--%" is displayed before a migration plan is made.
  - **Total:** Total number of logical volumes in the parity group.
  - **Res.:** Number of reserved volumes in the parity group.

Table 4.1 HDD Types versus Performance

Order	HDD Type	Capacity	
0	DKR2B	18 GB	*High Performance*
1	DKR1C	72 GB	
2	DKR1B	47 GB	*Low Performance*

Note: For the latest information on available HDD types, please contact your Hitachi Data Systems account team.

When you use the same maximum disk usage rate for all HDD classes, the performance of the HDD becomes the only factor used in determining auto migration plans. When you specify different usage limits for the HDD classes, you can bias the auto migration function to favor (or avoid) certain HDD types. The migration of high-usage volumes to higher HDD classes is expected to significantly improve host access to the volumes, which in itself can also have a large effect on subsystem performance. You should always perform careful analysis of monitor data collected after migration to verify the results of migration operations. The results of auto migration operations can help you to determine the appropriate disk usage limits for your 9900 configuration and operational environment.

To set the maximum disk usage for an HDD class:

1. On the HIHSM Main panel, select the **Auto…** button to open the Auto Migration panel (refer to Figure 4.12).
2. On the Auto Migration panel, select the **Class…** button to open the Class List panel (refer to Figure 4.13).
3. On the Class List panel, select the desired HDD class (e.g., Class A), and then select the **Usage…** button to open the Class Usage panel (refer to Figure 4.15).
4. On the Class Usage panel, enter the desired value in the **Maximum disk utilization** field, and then select **OK**. You are returned to the Class List panel.

4.7.3 Setting the Auto Migration Parameters

HIHSM provides full user-control over the auto migration function. The Execution Parameter Settings panel (see Figure 4.16) allows you to turn ON/OFF the auto migration function and configure the detailed auto migration parameters. To open the Execution Parameter Settings panel, select **Execution…** on the Auto Migration panel.

Caution: Make sure to schedule your HIHSM auto migration operations around the 9900 SVP auto-reboot time. When the SVP reboots, HIHSM auto migration activities are discontinued and will not resume until the next scheduled start time.

Caution: When an error condition exists in the 9900 subsystem, resource usage can increase or become unbalanced. Do not use 9900 monitor data collected during an error condition as the basis for making HIHSM auto migration plans.

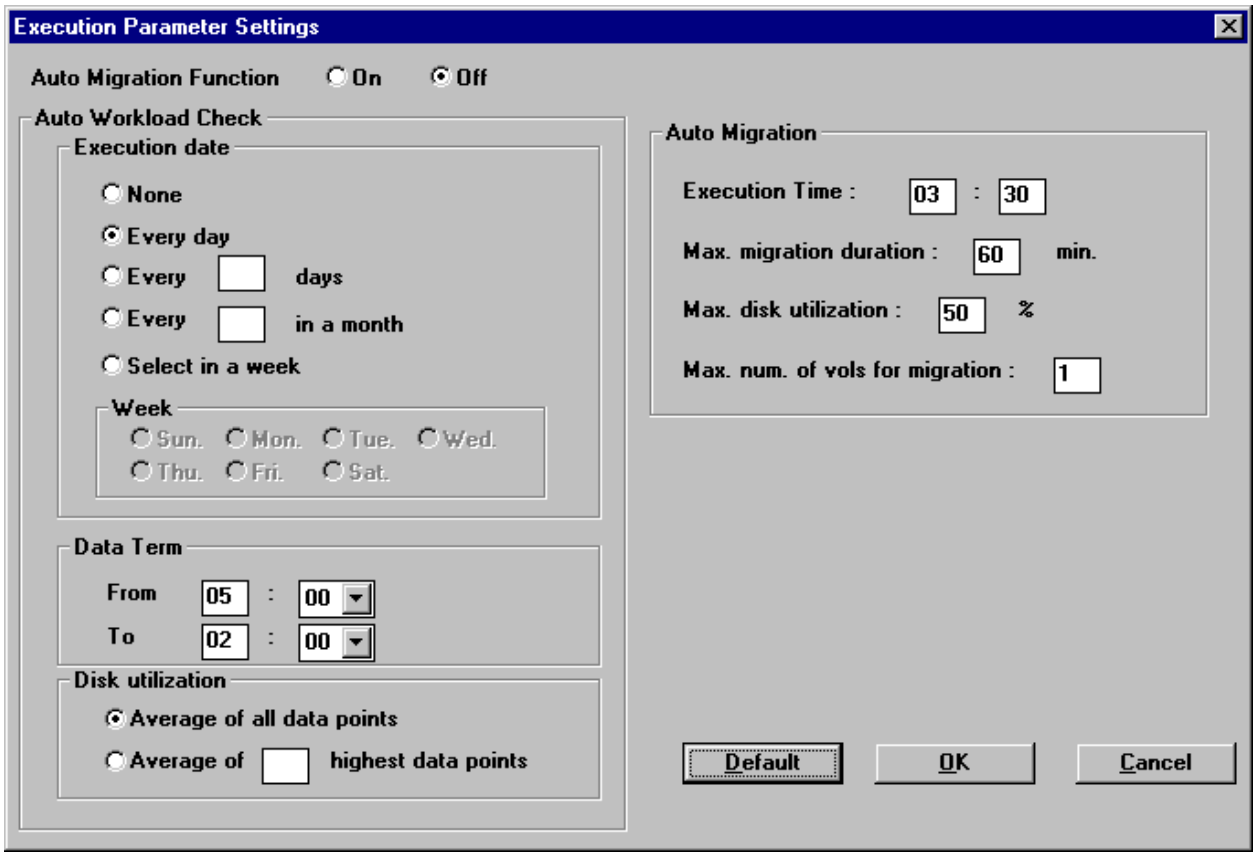

Figure 4.16 Execution Parameter Settings Panel

The **Auto Migration Function** buttons allow you to turn **On**/**Off** the auto migration function. The **Auto Workload Check** box allows you to set the parameters for making auto migration plans: execution date, data term to be analyzed, disk usage data samples to be analyzed. The **Auto Migration** box allows you to set the parameters for performing auto migration operations: execution time, maximum migration time, maximum disk utilization during auto migration, and maximum number of volumes which can be migrated at the same time.

The **Auto Workload Check** box allows you to set the following parameters for making auto migration plans. *Caution:* You should switch the monitoring ON before using this function.

- **Execution Date:**

None	Do not check the monitor data. HIHSM will not make new auto migration plans but will execute existing auto migration plans.
Every day	Make a plan every day.
Every X days	Make a plan every X days (X = 2-31).
Every X in a month	Make a plan on the Xth day of the month (X = 1-31). *Note:* If you specify 31, the last day of a month will be assumed for months having fewer than 31 days.
Select in a week	Make a plan on the selected day of the week.

- **Data Term:**

  From Start time of monitor data term to be analyzed: 00-23 hours, 15-min. intervals.

  To End time of monitor data term to be analyzed: 00-23 hours, 15-min. intervals.

  Important: Make sure that the data term and the auto migration execution time are separated by at least an hour. For example, if the data term is from 10:00 to 14:00, auto migration should not be performed between 09:00 and 15:00.

- **Disk Utilization:**

Average of all data points	Use the average usage rate for all monitor data samples to determine disk usage.
Average of X highest data points	Use the average usage rate for the X highest monitor data samples to determine disk usage.

The **Auto Migration** box allows you to set the following parameters for performing auto migration operations:

- **Execution Time:** Time to start auto migration operations on every selected day. *Important:* The following three restrictions apply to the **Execution Time**:
  - The **Execution Time** should be at least 15 minutes plus the time specified for **Max. migration duration** earlier than any automatic reboot or automatic data gathering.
  - The **Execution Time** should be at least 30 minutes later than any automatic reboot or automatic data gathering.
  - The Execution Time should be at least 15 minutes plus the time specified for **Max. migration duration** earlier than the start time specified for **Data Term**, or at least 60 minutes later than the end time specified for **Data Term**.

  Make sure that you specify the **Data Term** and **Execution Time** correctly, otherwise the auto migration will not be executed normally.

- **Max. migration duration:** Time limit (10-120 minutes) for auto migrations. If auto migration is not complete within this limit, HIHSM cancels the incomplete operations.
- **Max. disk utilization:** Disk usage limit (10-100%) during auto migration. If the disk usage for the source or target parity groups is over this limit during auto migration operations, HIHSM cancels and discards the auto migration plan.
- **Max. num of vols for migration:** Maximum number of volumes (1-40) that can be auto migrated at the same time.

To set the auto migration parameters:

1. Make sure that you have reserved the desired target volumes for HIHSM auto migration operations (see section 4.6).
2. On the HIHSM Main panel, select the **Auto…** button to open the Auto Migration panel.
3. On the Auto Migration panel, select the **Execution…** button to open the Execution Parameter Settings panel.
4. On the Execution Parameter Settings panel, make sure that the auto migration function is on. If not, turn the auto migration function **On** before continuing.
5. Select/enter the desired auto migration parameters in the **Auto Workload Check** box:

   Note: The default auto migration parameters are rough estimates only and may not provide the desired results. Please specify the appropriate values for your operational environment.

   a) Set the schedule for checking monitor data and making auto migration plans: none, every day, every X days, every Xth day of the month, or once per week on the specified day. If you select every X days, enter the desired number of days (e.g., 3 = check every three days).

   b) Enter the monitor data term that the auto migration function will analyze (**From** = start time of monitor data term, **To** = end time of monitor data term).

   c) Select the desired method for determining the average disk usage: **Average of all data points**, or **Average of X highest data points** (enter the desired number for X).
6. Enter the desired auto migration parameters in the **Auto Migration** box:

   a) Enter the desired start time for the scheduled auto migration plan. HIHSM will execute the existing auto migration plan (if any) at this time every day.

   b) Enter the desired maximum migration duration time (10-120 minutes).

   c) Enter the maximum disk usage rate (10-100%) during auto migration operations. For example, if you enter 60% and the disk usage rate during auto migration is higher than 60%, HIHSM will cancel the auto migration operation.

   d) Enter the desired maximum number of volumes that can be migrated at the same time by the auto migration function (1-40 volumes).
7. After confirming the auto migration parameters that you have selected, select **OK**.
   Reminder: Do not perform any manual migration operations while auto migration is on.

Note: If the 9900 SVP reboots during auto migration operations, HIHSM displays SVP Reboot on the Auto Migration History panel but does not record the migration results. Please refer to the main history log (see section 4.9) for the migration operation results. The interrupted auto migration operations are displayed as **Migration Status-Under Migration** on the Auto Migration panel. Please cancel those auto migration operations.

4.7.4 Making a New Auto Migration Plan

The **Make New Plan** button on the Auto Migration panel allows you to make a new auto migration plan at any time. To make a new auto migration plan:

1. On the HIHSM Main panel, select the **Auto…** button to open the Auto Migration panel.
2. Review the existing auto migration plan. If you want HIHSM to make a new auto migration plan, select the **Make New Plan** button. The existing plan will be overwritten by the new plan.
3. HIHSM deletes the old plan and displays the new plan in the **Generated Migration Plan** box. The **Next Migration Time** is the time that the new plan will be executed.

   Note: Under certain conditions HIHSM is not able make a new plan (e.g., no monitor data). If this occurs, select the **History Log…** button, and check the auto migration history log (see section 4.7.6) for detailed information on the failed make-new-plan operation.
4. Verify the new plan. If you do not agree with this plan, select the **Cancel Plan** button to remove this plan. If you select the **Make New Plan** button again, the plan will also be overwritten.

4.7.5 Canceling an Auto Migration Plan

The **Cancel Plan** button on the Auto Migration panel allows you to cancel the existing auto migration plan. To cancel an auto migration plan:

1. On the HIHSM Main panel, select the **Auto…** button to open the Auto Migration panel.
2. Review the existing migration plan on the Auto Migration panel, and confirm that you want to cancel the entire plan. If so, select the **Cancel Plan** button.
3. When the cancel plan confirmation message is displayed, select **OK** to cancel the plan, or select **Cancel** to cancel your request to cancel the migration plan.
4. If one or more auto migration operations are not removed from the **Generated Migration Plan** box, the auto migration plan is already in the process of being executed. You can use the **Delete** button on the Volume List panel (see section 4.8.3) to stop and cancel an auto migration operation which is in progress.

Note: Auto migration is usually planned for migrating multiple LDEVs in a consecutive sequence, and one LDEV after another should be migrated according to the sequence of those consecutive plans. When you make several auto migration plans, the sequence of migration may be important and the plans are likely to effect one another. Therefore, you are not allowed to delete one of the consecutive plans. If you delete a plan, the remaining plans specified to be executed after the deleted plan will also be deleted.

4.7.6 Auto Migration History Log

The HIHSM auto migration function logs all auto migration events in the auto migration history log. An auto migration event is any event related to the creation and execution of auto migration plans. The Auto Migration History Display panel (see Figure 4.17) displays the auto migration history log and allows you to delete all events from this log. To open the Auto Migration History Display panel, select the **Auto…** button on the HIHSM Main panel, and then select the **History Log…** button on the Auto Migration panel.

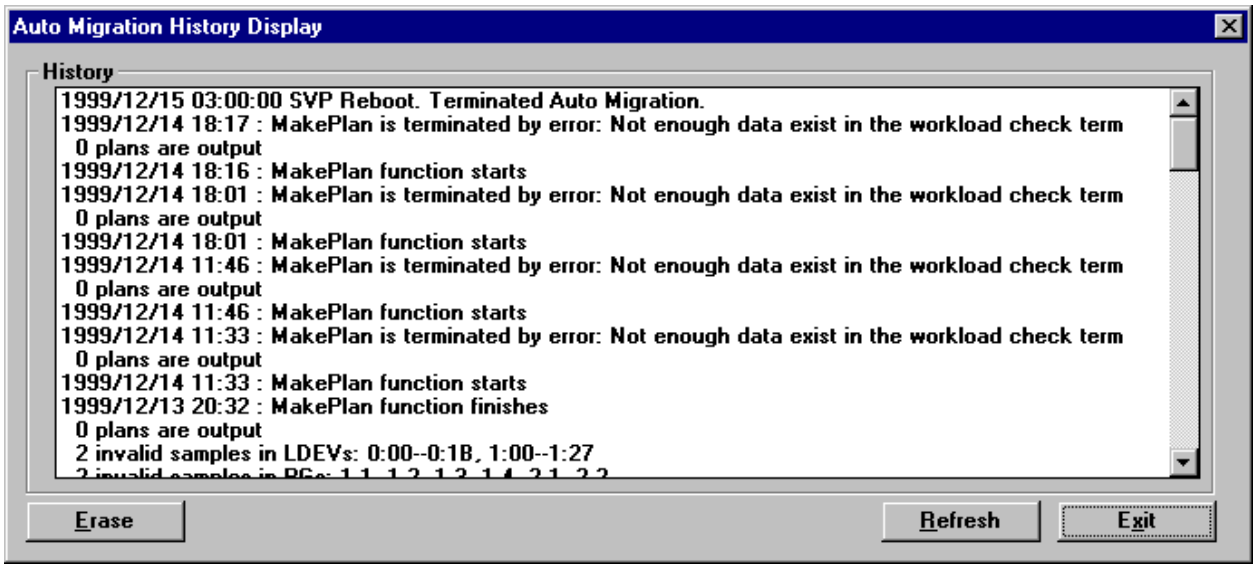

Figure 4.17 Viewing the Auto Migration History Log

The **Erase** button allows you (administrator only) to remove all entries from this history log. The **Refresh** button refreshes the information displayed on the panel. Table 4.2 lists all of the log entries that can be displayed on the Auto Migration History Display panel.

Note: The auto migration history function records all auto migration events. The HIHSM Main History panel (see section 4.9) displays a complete record of all HIHSM migration operations.

Table 4.2 Auto Migration History Log Entries (continues on the next page)

Normal end of migration:

Migration Complete (CU:LDEV->CU:LDEV) Start:yyyy/mm/dd hh:min:sec -> End:yyyy/mm/dd hh:min:sec
(a) (b) (c) (d)

(a) Source volume
(b) Target volume
(c) Copy start time
(d) Copy end time

Migration was canceled (over the limit duration):

Migration Canceled (CU:LDEV->CU:LDEV) Start:yyyy/mm/dd hh:min:sec -> End:yyyy/mm/dd hh:min:sec
(a) (b) (c) (d)

(a) Source volume
(b) Target volume
(c) Copy start time
(d) Copy end time

Migration was canceled (invalid parity group):

Migration Canceled (CU:LDEV (X-X)->CU:LDEV (X-X)) yyyy/mm/dd hh:min:sec (Invalid Parity Group)
(a) (b) (c) (d) (e)

(a) Source volume
(b) Parity group of the source volume
(c) Target volume
(d) Parity group of the target volume
(e) Canceled time

Migration was stopped (other reason):

Migration Stopped (CU:LDEV->CU:LDEV) yyyy/mm/dd hh:min:sec (XXXXXXXXXXXX)
(a) (b) (c) (d)

(a) Source volume
(b) Target volume
(c) Canceled time
(d) Reason for cancellation:

No reserve volume.	No reserved volume is set. Please make another plan.
Reserve volume emulation is different.	The reserved volume's emulation type is different.
Utilization check.	The usage rate is over the limit, or there is no monitor data.
Migration failed. Error code: XXXX.	The migration operation failed.
Reserve volume size is different.	Reserve volume size is different.
Reserve volume emulation is not supported.	The emulation type of reserved volume is not supported.
Utilization check failed.	The check of the usage rate finished abnormally.
Reserve volume check failed.	The check of the reserved volume finished abnormally.

Migration plan was deleted (because the previous plan had been deleted):

Migration Plan deleted (CU:LDEV->CU:LDEV) yyyy/mm/dd hh:min:sec (Pre-Plan is deleted)
(a) (b) (c)

(a) Source volume
(b) Target volume
(c) Canceled time

Started making an auto migration plan:

yyyy/mm/dd hh:min : MakePlan function starts
(a)

(a) Start time

Table 4.2 Auto Migration History Log Entries (continued)

Finished making an auto migration plan:

yyyy/mm/dd hh:min : MakePlan function finishes
(a)

(a) End time

Output the auto migration plans:

X plans are output
(a)

(a) Number of plans

Contents of auto migration plan:

PlanXXX: Src LDEV CU:LDEV in Grp X-X, Dst LDEV CU:LDEV in Grp X-X
(a) (b) (c) (d) (e)

(a) Plan number
(b) Source volume
(c) Parity group of the source volume
(d) Target volume
(e) Parity group of the target volume

Failed making an auto migration plan:

yyyy/mm/dd hh:min : MakePlan is terminated by error: XXXXXXXXXX
(a) (b)

(a) End time
(b) Reason for failure:

Cannot make proper migration plan.	The program could not make a proper migration plan. Check the number of the reserved volumes and their locations. Check the maximum disk usage and change if necessary.
Not enough HIHSM data in the monitor data term.	The program did not have enough data to make the plan in the target term. Check the monitored data and/or gather more monitor data. Check the monitor data term.
Failed to get XXXXXXXXXX (information or data).	The program failed to get the information or data to make a plan. In the case of information, make the initial value again by selecting [**Initialize**]. In the case of data, check the status of gathered data and get them again.
Failed to write to XXXXXXXXXX (information or data).	The program failed to write the information to make a plan.
Invalid XXXXXXXXXX (information or data).	The required information/data could not be used because the information was invalid. In the case of information, make the initial value again by selecting [**Initialize**]. In the case of data, check the status of gathered data and get them again.
Memory allocation error.	The program failed to allocate the memory to make the plan.

Log entry that the auto migration plan could not be made:

Cannot make plan: Class X Grp X-X
(a) (b)

(a) Class
(b) Parity group

Some monitor data samples were invalidated:

Grp X-X: X samples are invalidated because of migration.

To reduce the influence of migration during the monitor term, HIHSM made a plan without using some data samples.

Table 4.2 Auto Migration History Log Entries (continued)

New volumes have been installed: New entries are added for following LDEVs: CU:LDEV, CU:LDEV, ...
Invalid data: Too many invalid data: invalidated all data The contents of data in the indicated group or volumes is invalid. HIHSM made a plan without those data.
Volume utilization check failed: Utilization check failed Check volume usage failed because the necessary data could not be obtained. Please get the data again.
Volume reserve attribute check failed: Reserve volume check failed Check volume reserve attribute failed because the necessary data could not be obtained. Please get the data again.

4.8 Starting, Checking, and Stopping Migration Operations

Note: Make sure to confirm the latest configuration information on the HIHSM Main panel before and after migration.

The Volume List panel (see Figure 4.18) allows you to start manual migration operations, check the status of in-process manual and auto migration operations, and stop in-process manual and auto migration operations. (An auto migration operation can only be started by the auto migration function.) To open the Volume List panel, select the desired parity group on the HIHSM Main panel, and then select the **Volume List…** button.

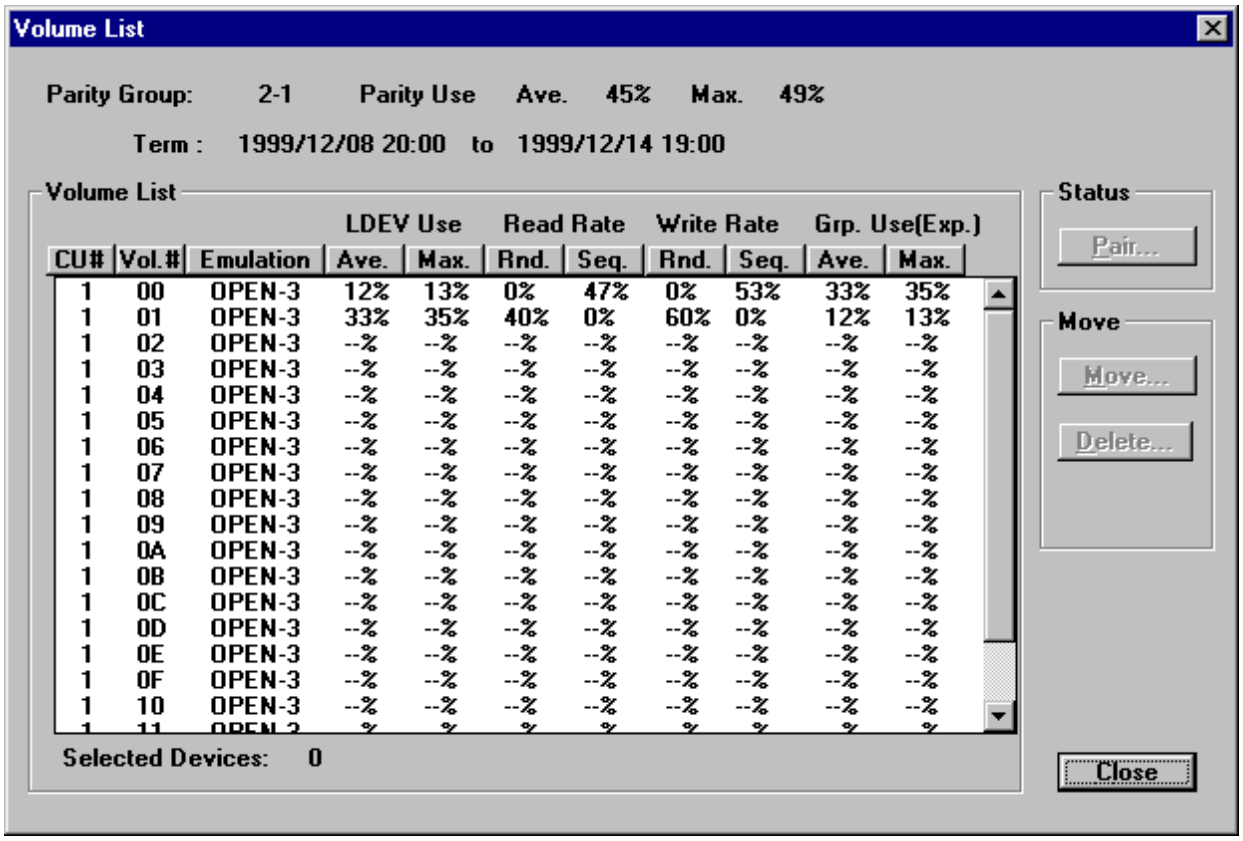

Figure 4.18 Starting, Stopping, and Checking Migration (Volume List Panel)

The **Pair…** button opens the Pair Status panel (see section 4.8.2), which displays the HIHSM pair information for the selected volume and the status of the HIHSM migration operation (manual or automatic).

The **Move…** button opens the Select Target Group panel (see section 4.8.1), which allows you to start a manual migration operation for the selected source volume.

The **Delete…** button opens the Delete Pair Confirmation panel (see section 4.8.3), which allows you to stop one or more in-process migration operations (manual or automatic).

4.8.1 Starting a Manual Migration Operation

After studying the desired HIHSM monitor data, you may want to perform one or more specific migration operations to address a specific performance issue. Before you can start a manual migration operation, you must:

- Identify the source volume (the volume that you want to migrate),
- Identify the target (destination) volume for the manual migration operation,
- Make sure that the desired target volume is reserved (refer to section 4.6), and

Caution: When an error condition exists in the 9900 subsystem, resource usage can increase or become unbalanced. Do not use 9900 monitor data collected during an error condition as the basis for planning HIHSM manual migration operations. Also, when you continuously perform manual migration operations, the configuration may be changed after each migration is complete and before the next consecutive operations will be performed. Therefore, make sure to confirm the parity group specified for migration destination before performing each migration operation.

When you are ready to start manual migration operations, you must turn off the auto migration function and cancel all existing auto migration plans (refer to section 4.7). If you do not turn off auto migration and cancel existing plans, the auto migration function may experience errors (e.g., target volume no longer available). In addition, the usage data on which the auto migration plan was based may no longer be valid after you perform manual migration operations.

The Select Target Group panel (see Figure 4.19) displays the selected HIHSM source volume and the available target parity groups and allows you to select the desired target parity group. The available target parity groups are the parity groups which contain reserved volumes. To open the Select Target Group panel, select the desired HIHSM source volume on the Volume List panel, and then select the **Move…** button.

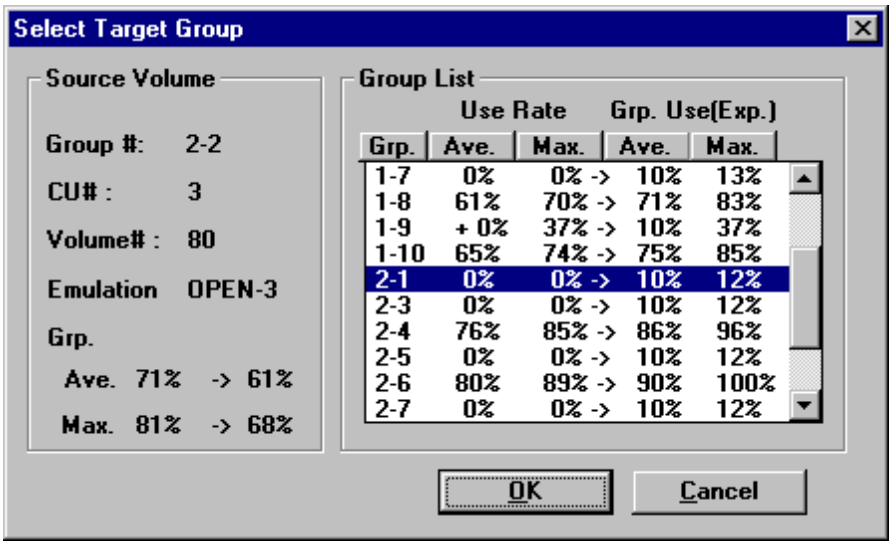

Figure 4.19 Select Target Group Panel

The **Source Volume** box displays the following information for the selected volume: parity group number, CU number, volume number, emulation type, and estimated group usage after the proposed migration. The **Group List** box displays the available target parity groups and their current usage rates (**Use Rate**) and estimated (expected) usage rates (**Grp. Use (Exp.)**). HIHSM uses the currently specified monitor data term to calculate estimated usage.

To start a manual migration operation:

1. On the HIHSM Main panel, select the parity group which contains the desired source volume, and then select the **Volume List…** button to open the Volume List panel.
2. On the Volume List panel, select the desired source volume, and then select the **Move…** button to open the Select Target Group panel.
3. After reviewing the expected usage rate information on the Select Target Group panel, select the desired target parity group, and then select **OK**.
4. The Select Target Volume panel (see Figure 4.20) now opens and displays the available target volumes (i.e., reserved volumes) in the selected target parity group. Select the desired target volume, and then select **OK**.
5. The Add Pair Confirmation panel (see Figure 4.21) displays the selected source and target volumes and the expected usage rates after migration. After reviewing this information, select **OK** to start the specified manual migration operation, or select **Cancel** to cancel your request to start migration.

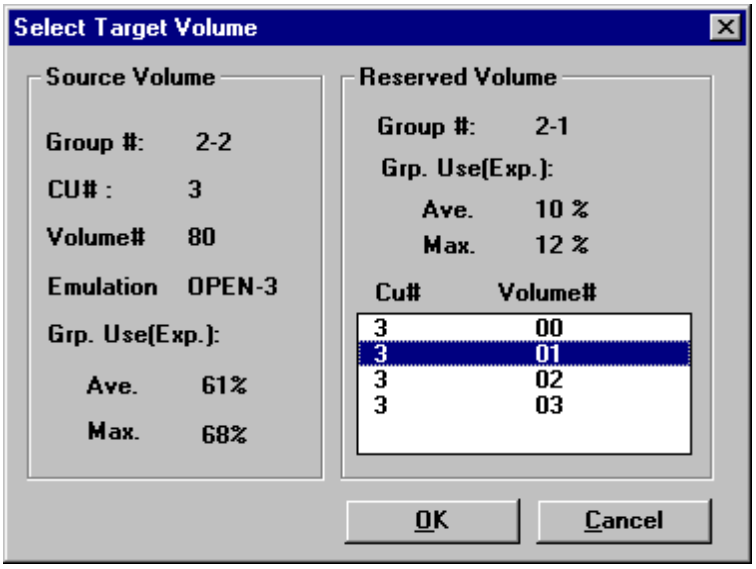

Figure 4.20 Selecting the Target Volume

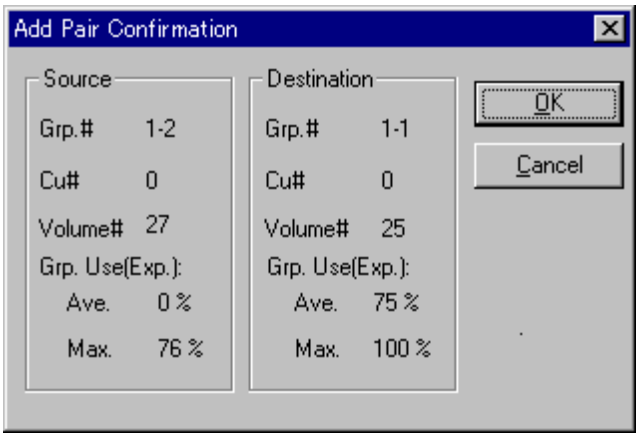

Figure 4.21 Add Pair Confirmation Panel

4.8.2 Viewing Migration Status

HIHSM allows you to view the status of a manual or automatic migration operation while the operation is in progress. The Pair Status panel (see Figure 4.22) displays the selected source volume, the destination volume in the HIHSM pair, and the status of the HIHSM migration operation. To open the Pair Status panel, select the desired source or target volume on the Volume List panel, and then select the **Pair…** button.

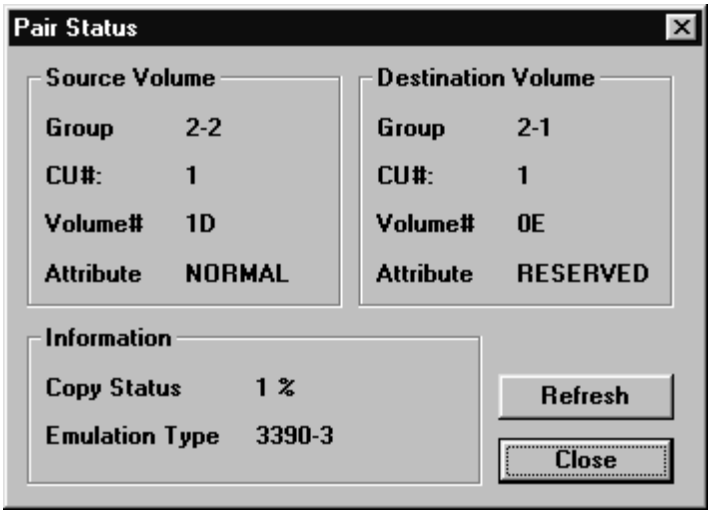

Figure 4.22 Pair Status Panel

The Pair Status panel displays the following information:

- **Source Volume**: parity group ID, CU image, LDEV ID, and reserve attribute.
- **Destination Volume**: parity group ID, CU image, LDEV ID, and reserve attribute.
- **Information**: copy status (percent completion of HIHSM migration operation), and device emulation type.

To view the status of an in-process HIHSM migration operation:

1. On the HIHSM Main panel, select the parity group which contains the desired volume, and then select the **Volume List…** button to open the Volume List panel.
2. Select the desired volume on the Volume List panel, and select the **Pair…** button to open the Pair Status panel. Use the **Refresh** button to update the copy status information displayed on the Pair Status panel.

4.8.3 Stopping a Migration Operation

If you need to stop/abort a HIHSM migration operation (manual or automatic), HIHSM allows you to an delete HIHSM pair while the migration is in process. You can delete a HIHSM pair any time before completion of the migration operation. ***Note:*** When a HIHSM pair is deleted during the migration operation, the data on the target volume is not guaranteed.

The Delete Pair Confirmation panel (see Figure 4.23) displays the requested HIHSM pair delete operation and allows you to delete the specified migration pair(s). To open the Delete Pair Confirmation panel, select the desired volume(s) on the Volume List panel, and then select the **Delete…** button.

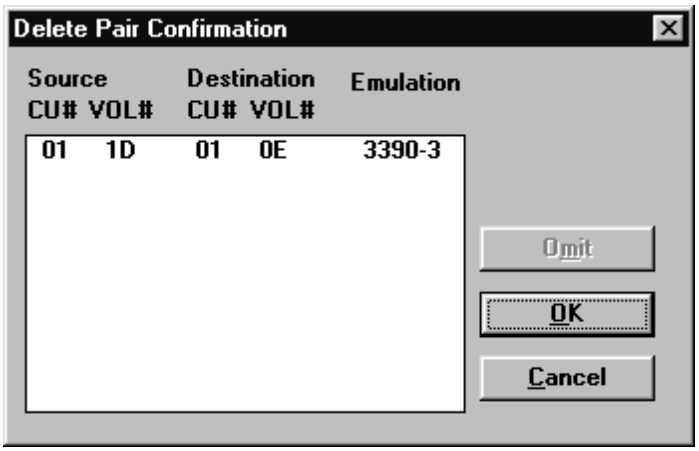

Figure 4.23 Stopping Manual Migration

The **Omit** button allows you to remove one or more HIHSM volume pairs from the list Delete Pair Confirmation panel. To stop one or more HIHSM migration operations:

1. On the HIHSM Main panel, select the parity group which contains the desired volume(s), and then select the **Volume List…** button to open the Volume List panel.
2. The Volume List panel displays an asterisk (*) by volumes are currently being migrated. On the Volume List panel, select the desired volume(s), and then select the **Delete…** button to open the Delete Pair Confirmation panel.
3. The Delete Pair Confirmation panel displays the selected volume pair(s). To remove a HIHSM volume pair from the list, select it and then select **Omit**. To cancel your request to delete the HIHSM pair(s), select **Cancel**.
4. After you have reviewed and confirmed the information displayed on the Delete Pair Confirmation panel, select **OK** to delete the specified HIHSM pair(s). The specified HIHSM migration operation(s) is/are aborted after you select **OK**.

4.9 Viewing the Main History Log

All HIHSM volume migration operations are logged in the main history log. The HIHSM main history log displays a record of each individual migration operation, whether automatic or user-requested. For information on auto migration events (e.g., creating auto migration plans), refer to the auto migration history log (see section 4.7.6).

The History Display panel (see Figure 4.24) displays the HIHSM main history log and allows you to delete all entries from this log. To open the History Display panel, select the **History…** button on the HIHSM Main panel.

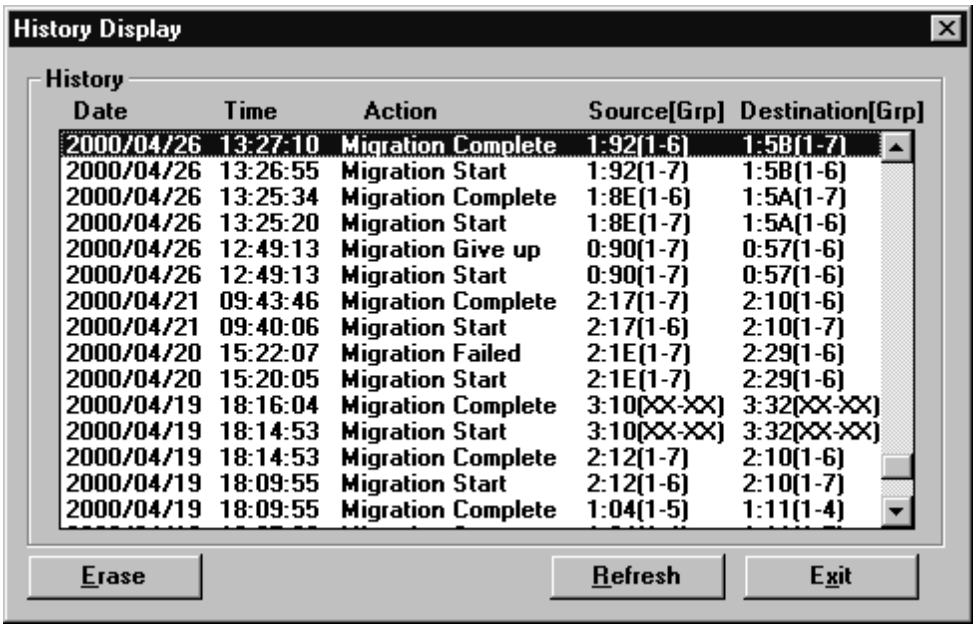

Figure 4.24 Viewing the Main History Log

The History Display panel displays the following information for each migration operation:

- **Date** of the specified migration action.
- **Time** of the specified migration action.
- **Action:**
  - **Migration Start**: The migration operation started.
  - **Migration Complete**: The migration operation completed successfully.
  - **Migration Cancel**: The migration operation was canceled by the user.
  - **Migration Failed**: The migration operation failed (R-SIM generated).
  - **Migration Give Up**: The migration operation was canceled by HIHSM (e.g., usage rates during operation exceeded user-specified maximum).

- **Source(Grp):** Source volume (LDEV ID) and parity group.
- **Destination(Grp):** Target volume and parity group.

  Note: Parity groups that have been removed by configuration changes are indicated: **xx-xx**.

The **Erase** button allows you (administrator only) to remove all entries from the history log. The **Refresh** button refreshes the information displayed on this panel. The **Exit** button closes this panel and returns you to the HIHSM Main panel.

Chapter 5 Installing and Using GraphTool

The GraphTool software for Dynamic Optimizer (HIHSM) enables you to extract the 9900 monitor data from the **artlog.lzh** export file (see section 4.4) and graph the data in several graphical formats. These graphical representations of 9900 monitor data enable you to identify workload characteristics and highlight key information about 9900 subsystem usage such as peaks and trends. The GraphTool software is included in the HIHSM software product.

This chapter describes the following procedures:

- Installing and uninstalling the GraphTool software (see section 5.1),
- Starting and exiting the GraphTool software (see section 5.2),
- Extracting the data from the log file (see section 5.4),
- Displaying the data graphs (see section 5.5),
- Viewing the data graphs (see section 5.6),
- Printing the data graphs (see section 5.7), and
- Closing the data graphs (see section 5.8).

5.1 Installing and Uninstalling GraphTool

The GraphTool software requires the following operational environment:

- A Windows®-based PC (Windows® 95, 98, 2000, or Windows NT® 4.0) with a diskette drive (to access the **a:\artlog.lzh** file) and a CD drive (to install the GraphTool software).
- At least eight MB of free space (three MB to install GraphTool, about five MB when using GraphTool). More free space may be required depending on the configuration of the 9900 subsystem whose monitor data you are analyzing (e.g., number of parity groups, DKPs, volumes) and the range of monitor data being analyzed (i.e., size of **artlog.lzh**).

5.1.1 Installing the Graphing Function

To install the HIHSM GraphTool software:

1. Run **setup.exe** in the installation CD-ROM. The Welcome panel opens.
2. On the Welcome panel, select **Next** to open the Choose Destination Location panel (see Figure 5.1). If you want to use the default installation folder, select **Next**. If you want to change the installation folder, select **Browse...** and enter the desired installation location.
3. Installation starts, and the files are copied to the specified folder (see Figure 5.2).
4. When installation is complete, the Setup Complete panel opens (see Figure 5.3).
5. Confirm that the files specified in Table 5.1 were created.

   Note: Certain Windows® folder and file display options may hide some of these files (e.g., the two **dll** files) and/or the file extensions. Please change the folder and file settings so that you can verify all HIHSM graphing function files by file name and file extension.

Table 5.1 Graphing Function Components

File Name	Function
GTool.exe	Graphing function execution.
GraphTool.exe	Graphing tool launcher displaying the copyright dialog box. (This file gives parameters for starting the graphing tool.)
MFC42.dll	Dynamic link library for MFC.
MSVCRT.dll	Dynamic link library for starting the graphing tool.
R40cmn.dll	Dynamic link library for 9900 Remote Console program.
RMTCMN32.dll	Dynamic link library for 7700E Remote Console program.
unlha32.dll	Dynamic link library for unpacking.
DATRD300.dll	Dynamic link library for reading 7700E graphing tool data.
DATRD400.dll	Dynamic link library for reading 9900 graphing tool data.
GRTDATRD.dll	Dynamic link library for reading graphing tool data.
oemtype	OEM type data

Figure 5.1 Choose Destination Location Panel

Figure 5.2 Installing GraphTool

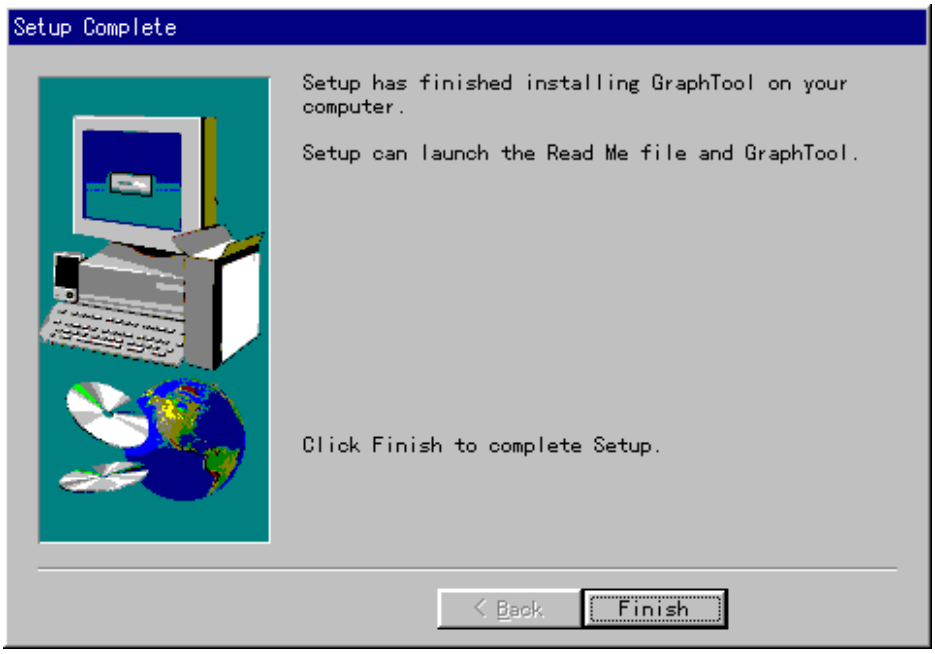

Figure 5.3 Setup Complete Panel

5.1.2 Viewing the GraphTool Version Information

To check the GraphTool program version information:

1. On the GraphTool main panel (see Figure 5.6), select **About GraphTool...** from the **Help** menu, or select the ? icon to display the GraphTool version information.
2. Select the **OK** button to close the version information window.

5.1.3 Uninstalling the GraphTool

To uninstall the HIHSM graphing function (GraphTool):

1. From the Windows **Start** menu, select **Setup**, select **Control Panel**, and then select **Adding and Deleting Applications**.
2. Select **GraphTool**, and then select the **Add and Delete** button.
3. When the confirmation message appears, select **Yes** to uninstall the GraphTool. The Uninstall Shield starts, and the corresponding program files are deleted (see Figure 5.4).

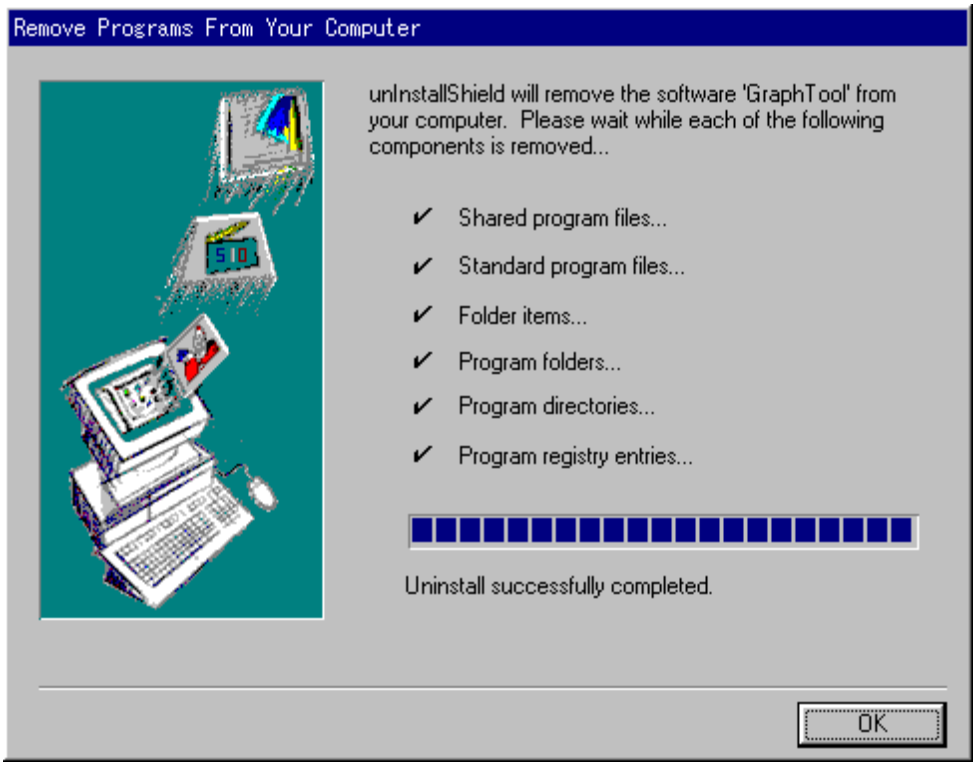

Figure 5.4 Uninstalling GraphTool

5.2 Starting and Exiting GraphTool

To start the GraphTool software:

1. From the Windows **Start** menu, select **Program Files**, and then select **GraphTool**. Or, you can double-click the **GraphTool.exe** file in the GraphTool installation directory.
2. The GraphTool Main panel opens (see section 5.3). The GraphTool Main panel provides access to all GraphTool functions.

To exit the GraphTool software:

1. On the GraphTool Main panel, select the **File** menu, and then select **Exit**.
2. If a GraphTool window is open, a confirmation message is displayed (see Figure 5.5). Select **Yes** to exit GraphTool, or select **No** to cancel your request.

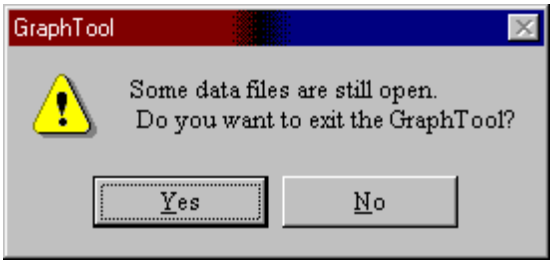

Figure 5.5 Exiting GraphTool

5.3 GraphTool Main Panel

When you start the GraphTool software, the GraphTool Main panel opens (see Figure 5.6). The GraphTool Main panel provides access to all GraphTool functions.

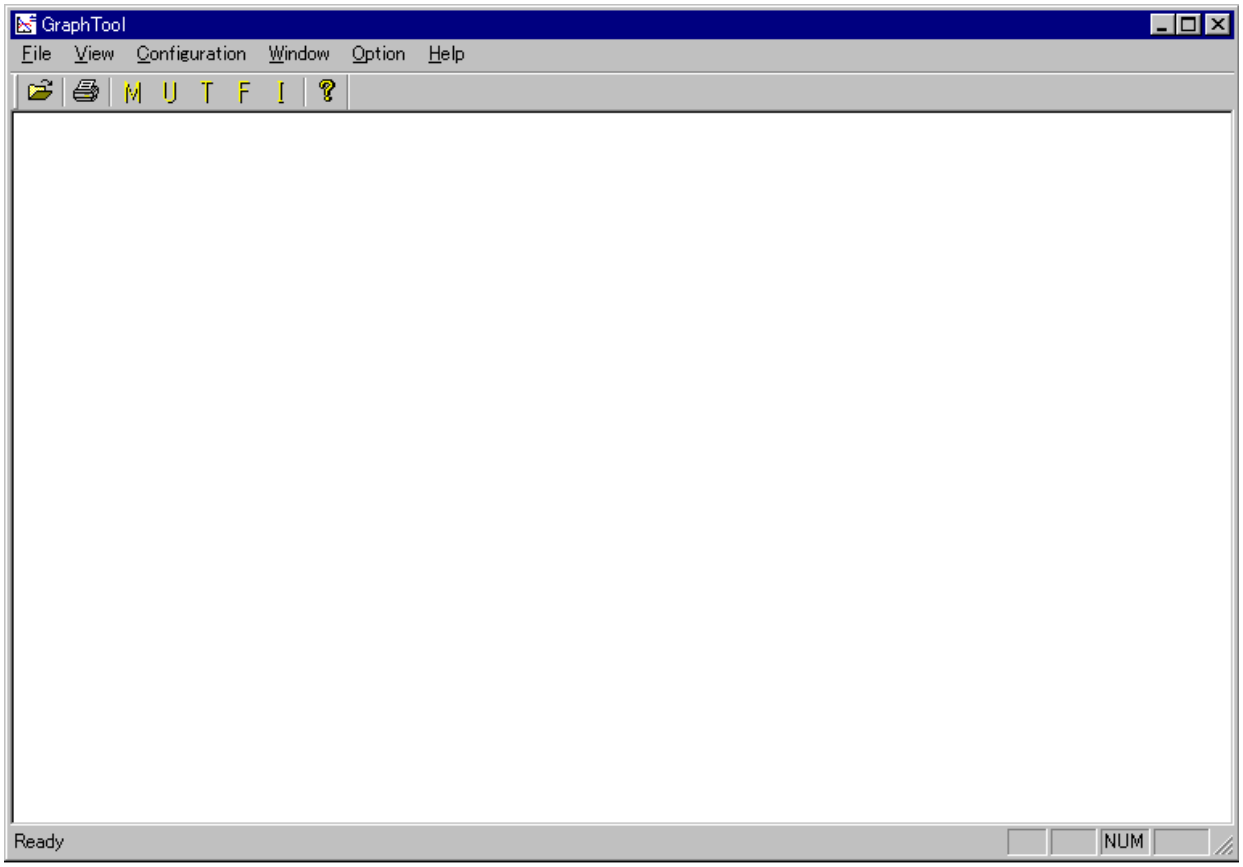

Figure 5.6 GraphTool Main Panel

The **File** menu provides the following commands for manipulating data files:

- **Open** ( ): Opens a data file for which you want to create a graph.
- **Close**: Closes the current data file which is open for graphing.
- **Print** ( ): Prints the graph currently displayed on the screen.
- **Print View**: Displays the print preview image for the current graph.
- **Print Setup**: Displays the screen for setting up the print options.
- **Recent File**: Opens the most recently edited and closed file.
- **Exit**: Terminates the GraphTool software, and closes the GraphTool Main panel.

The **View** menu displays the following commands for specifying the window features:

- **Tool Bar**: Allows you to specify whether or not to display the toolbar.
- **Status Bar**: Allows you to specify whether or not to display the status bar.

The **Configuration** menu provides the following commands for creating graphs:

- **Module** (): Opens the Select Module panel, which allows you to specify the graph module to be displayed (see section 0).
- **Unit** (): Opens the Select Unit panel, which allows you to select specific LDEVs to be displayed on the graphs (see section 5.5.4).
- **Term** (): Opens the Select Term panel, which allows you to specify the data term to be displayed on the graphs (see section 5.5.3).
- **Font** (): Opens the Font panel, which allows you to change the font settings for the graphs (see section 5.6.4).
- **Show Data Information** (): Opens the screen which displays the log file information.

The **Window** menu provides the standard commands for displaying the open window(s):

- **New Window**: Opens a new graph display window.
- **Cascade**: Displays multiple graph display screens (windows) on different layers.
- **Tile**: Displays multiple graph display screens (windows) by arranging them on a screen.
- **Arrange Icons**: Displays multiple graph display screens (windows) as icons arranged on the bottom left side of the GraphTool Main panel.

The **Option** menu provides the following commands for setting options:

- **Default Directory**: Allows you to change the default log data directory.
- **Save Option**: Allows you to save your changes to the default log data directory.

The **Help** menu provides the **About GraphTool** command (), which displays the version information for the GraphTool software.

The toolbar provides the following icons for the frequently used menu commands:

-  : **File-Open** command.
-  : **File-Print** command.
-  : **Configuration-Module** command (opens the Select Module panel).
-  : **Configuration-Unit** command (opens the Select Unit panel).
-  : **Configuration-Term** command (opens the Select Term panel).
-  : **Configuration-Font** command (opens the Font panel).
-  : **Configuration-Show Data Information** command.
-  : **Help-About GraphTool** command.

5.4 Extracting the Data from the Log File

The Read Data panel (see Figure 5.7) allows you to specify the location of and extract the data from a HIHSM export file (log data). After you have extracted the data from the specified file, you can start displaying the data as graphs (see section 5.5). To open the Read Data panel, select **Open** from the **File** menu or select the  icon on the GraphTool Main panel.

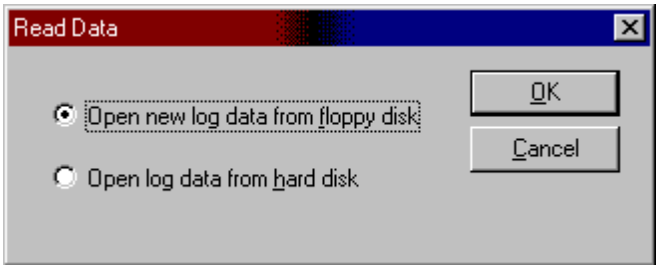

Figure 5.7 Read Data Panel

The **Open new log data from floppy disk** option allows you to read and extract data stored on one or more floppy disks (see section 5.4.1). The **Open log data from hard disk** option allows you to read and extract data stored on the hard disk drive (see section 5.4.2). When reading data from floppy disk, GraphTool also saves the data on the hard disk drive.

WARNING: Only one statistical information file (log data) can be stored in the log data directory. The user is responsible for backing up log data as needed prior to reading and extracting data from floppy disk.

5.4.1 Extracting Log Data on Floppy Disk

To extract a log data file stored on one or more floppy disks:

1. Open the Read Data panel (**File-Open**) (refer to Figure 5.7), select **Open new log data from floppy disk**, and then select the **OK** button.
2. The Log Data Directory panel opens (see Figure 5.8). On the Log Data Directory panel, specify the GraphTool log data directory (folder). GraphTool will save the log data on the floppy disk(s) into this directory. To find an existing directory, select the **Browse** button, and then locate and select the directory. Select the **Next** button to continue:
   a) If the specified directory does not exist, GraphTool displays a confirmation message for creating the new directory (see Figure 5.9). Select **OK** to create the new directory, or select **Cancel** to return to the Log Data Directory panel.
   b) If a log data file (**log.lzh**) already exists in the specified directory, GraphTool displays a confirmation message. If you want to *OVERWRITE* the existing log data with the data stored on floppy disk, select **OK**. If not, select **Cancel** to return to the Log Data Directory panel.
3. When the log data directory is ready, the Number of Floppy Disks panel opens (see Figure 5.10). Enter the number of floppy disks containing the statistical data, and select **Next**.
4. The Insert Floppy Disks panel opens (see Figure 5.11) and displays the disk number to be inserted (e.g., #1/1, #1/3). Insert the specified floppy disk into the disk drive, and select **OK**. To return to the Number of Floppy Disks panel select **Cancel**. Repeat this step as needed until the data has been read from all specified floppy disks.
5. When the data has been read from all floppy disks, a message prompts you to eject the last floppy disk (see Figure 5.12). Eject the floppy disk from the drive, and select **OK**.
6. The Extract Data panel now opens (see Figure 5.13), and data extraction begins in the specified folder on your hard drive (from step 2).
7. When data extraction is complete, GraphTool loads the initial data and displays the following information about the data file (see Figure 5.14): subsystem model (RAID400 = 9900), subsystem serial number, and data term (range of available data).

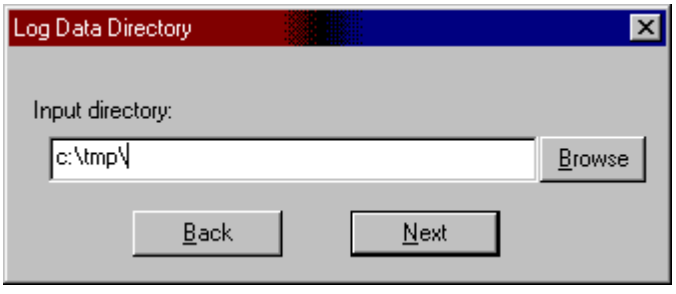

Figure 5.8 Log Data Directory Panel

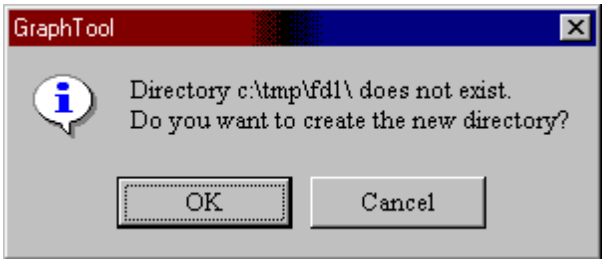

Figure 5.9 Creating a New Log Data Directory

Figure 5.10 Number of Floppy Disks Panel

Figure 5.11 Insert Floppy Disk Panel

Figure 5.12 Eject Floppy Disk Message

Figure 5.13 Extract Data Panel

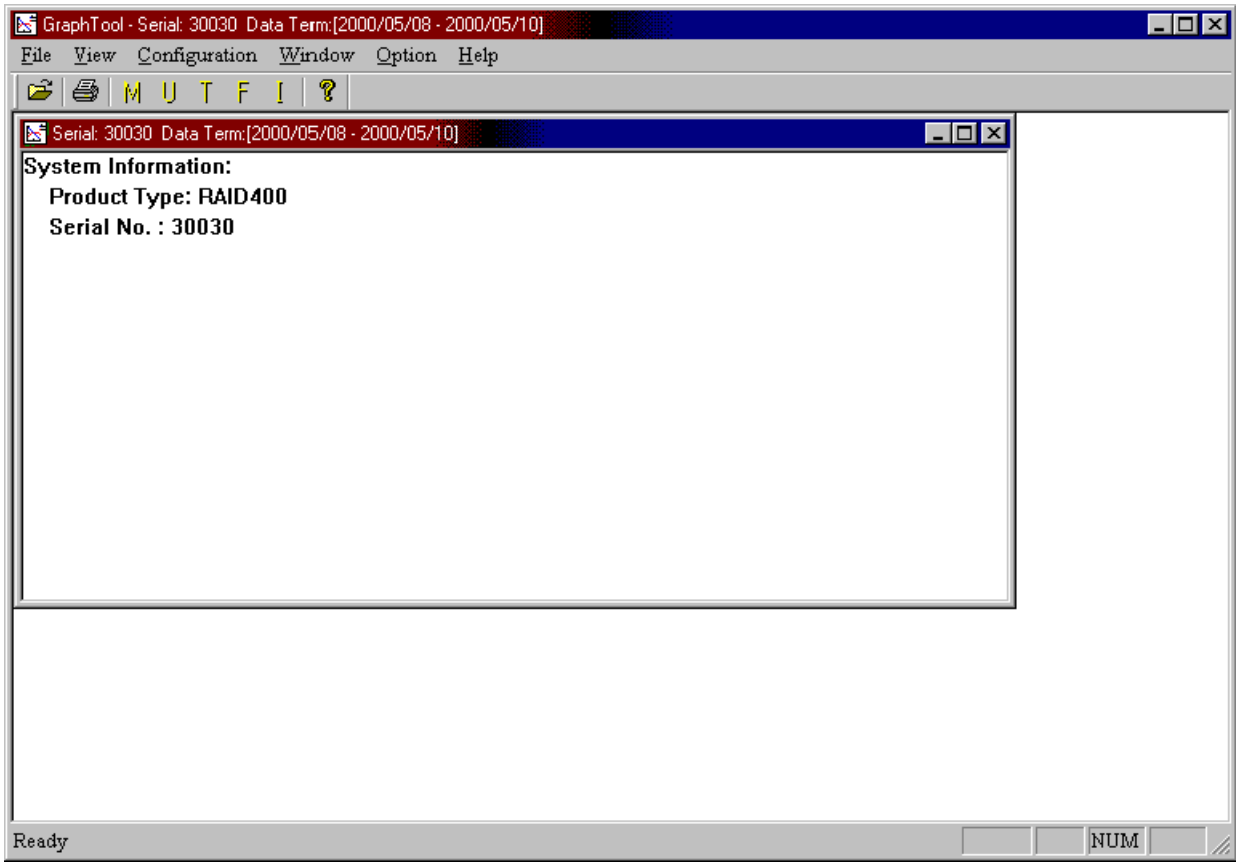

Figure 5.14 Data File Information

5.4.2 Extracting Log Data on the Hard Disk

To extract a log data file stored on the hard disk drive:

1. Open the Read Data panel (**File-Open**) (refer to Figure 5.7), select **Open log data from hard disk**, and then select the **OK** button.
2. When prompted, specify the desired log data file (e.g., **log.lzh**), and select the **Open** button. To open the last file you previously opened, select **Recent File** from the **File** menu.
3. The Extract Data panel opens (refer to Figure 5.13), and data extraction begins.
4. When data extraction is complete, GraphTool loads the initial data and displays the following information about the data file (refer to Figure 5.14): subsystem model (RAID400 = 9900), subsystem serial number, and data term (range of available data).

5.4.3 Specifying the Default Log Data Directory

As a program option, GraphTool allows you to specify a directory (folder) as the default log data directory. This location will be entered automatically during GraphTool operations, so that you do not have to enter the log data directory location.

To specify the default log data directory for GraphTool operations:

1. On the GraphTool Main panel, select the **Option** menu, and then select **Default Directory** to open the Option panel (see Figure 5.15).
2. On the Option panel, enter the desired directory in the **Default log data directory** text box using an absolute path (e.g., c:\tmp\logdata\). To find an existing directory, select the **Browse** button, locate the desired directory, and select **OK**.
3. When the directory entered on the Option panel is correct, select **OK**. If the specified directory does not exist, a confirmation message for creating the new directory is displayed. Select **OK** to create the new directory, or select **Cancel** to return to the Option panel.
4. After you have set the new default log data directory location, you need to save this information. Select the **Option** menu on the GraphTool Main panel, and then select **Save Options** to save the new log data directory location.
5. When the save completion message appears (see Figure 5.16), select **OK**.

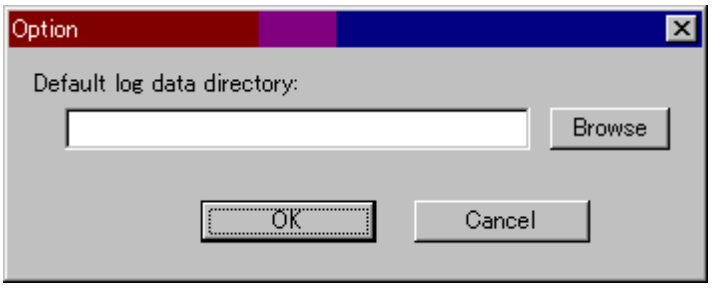

Figure 5.15 Option Panel

Figure 5.16 Completion Message

5.5 Displaying the Graphs

GraphTool allows you to specify the following parameters when displaying data graphs:

- Data module (e.g., CHP usage, DKP usage) (see section 5.5.1),
- Parity group (for the **Disk Utility (LDEV)** module only) (see section 5.5.2),
- Data term (start and end date/time), plotting scale (data interval), and plotting method (actual, average, or maximum values) (see section 5.5.3), and
- LDEV(s) (for the **Disk Utility (LDEV)** module only) (see section 5.5.4).

Note: While graphs are being displayed, you cannot perform any other GraphTool operations.

To display the 9900 monitor data as graphs:

1. Make sure that the desired log data file has been extracted using the Read Data panel (refer to section 5.4).
2. Open the Select Module panel (see Figure 5.17), select the data module that you want to display, and then select **OK**. If you select the **Disk Utility (LDEV)** module, the Select Parity Group panel opens automatically (see Figure 5.18). Select the desired parity group, and select **OK**.
3. After the data module has been selected, the Select Term panel opens automatically (see Figure 5.19). Select the desired data term, step (graph scale), and plot mode, and then select **OK**. The monitor data is loaded from the location specified on the Read Data panel.
4. After the data term has been selected and the data has been loaded, the Select Unit panel opens automatically (see Figure 5.20). Select the desired units (LDEVs) for which you want to display graphs. If you want to check a graph for an LDEV before continuing, select the **Preview** button. When you are finished selecting LDEVs and previewing graphs, select **OK** to close the Select Unit panel and plot the specified data graphs.

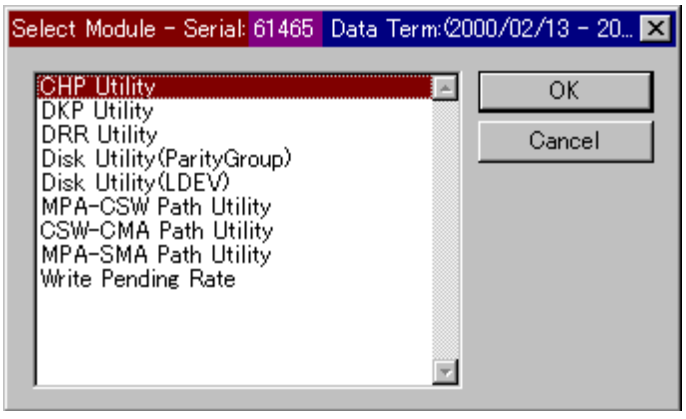

Figure 5.17 Select Module Panel

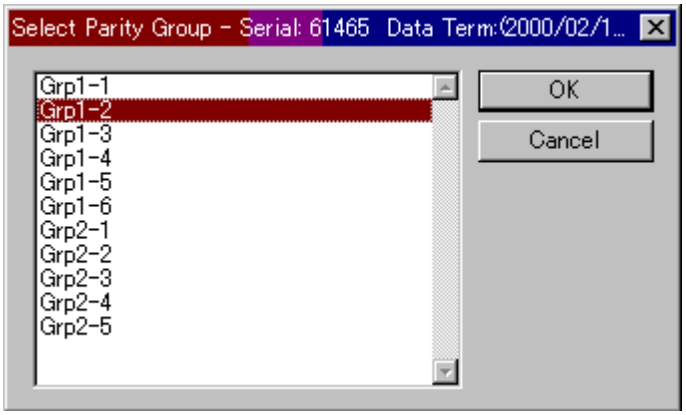

Figure 5.18 Select Parity Group Panel

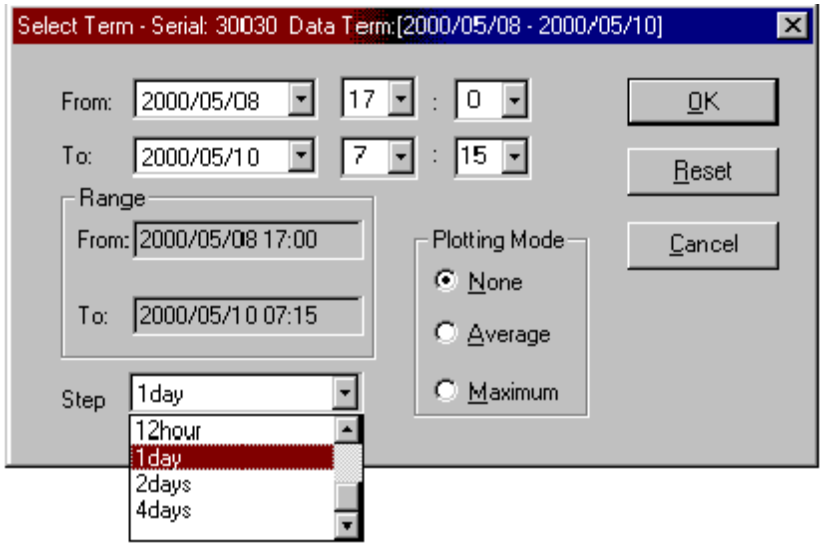

Figure 5.19 Select Term Panel

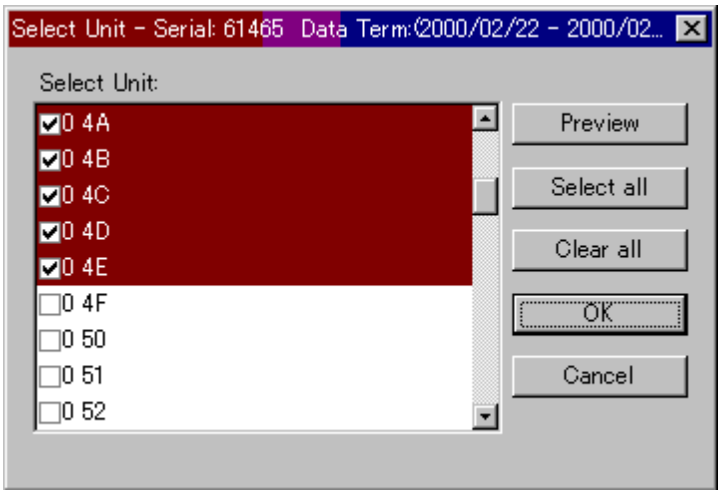

Figure 5.20 Select Unit Panel

5.5.1 Selecting the Data Module

The Select Module panel (see Figure 5.21) allows you to specify the data module to display. To open the Select Module panel, select **Module** under the **Configuration** menu, or select the  icon on the toolbar on the GraphTool panel.

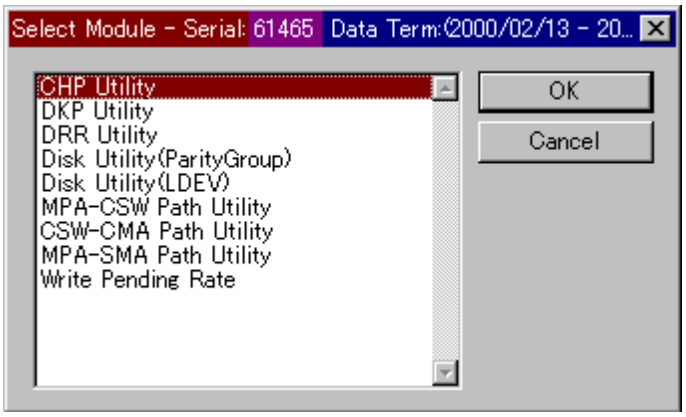

Figure 5.21 Select Module Panel

The **Module** list box lists the data modules available for graphing:

- **CHP Utility**: Actual, average, and maximum usage rates of the CHPs.
- **DKP Utility**: Actual, average, and maximum usage rates of the DKPs.
- **DRR Utility**: Actual, average, and maximum usage rates of the DRRs.
- **Disk Utility (Parity Group)**: Actual, average, and maximum usage rates of all parity groups.
- **Disk Utility (LDEV)**: Actual, average, and maximum usage rates of LDEVs of a specific parity group. If you select this module, the Select Parity Group panel opens automatically to allows you to specify the desired parity group (see section 5.5.2).
- **MPA-CSW Path Utility**: Actual, average, and maximum usage rates of the paths between the adapters and CSWs (Adapter-CSW).
- **CSW-CMA Path Utility**: Actual, average, and maximum usage rates of the paths between the CSWs and cache (CSW-Cache).
- **MPA-SMA Path Utility**: Actual, average, and maximum usage rates of the paths between the adapters and SMs (Adapter-SM).
- **Write Pending Rate**: Actual, average, and maximum rates of queued write I/Os.

If you selected the **Disk Utility (LDEV)** module, the **OK** button opens the Select Parity Group panel (see section 5.5.3) to allow you to select the desired parity group. If you did not select the **Disk Utility (LDEV)** module, the **OK** button opens the Select Term panel (see section 5.5.3) to allow you to specify the range of data to be displayed. The **Cancel** button cancels your request and closes the Select Module panel.

5.5.2 Selecting the Parity Group

The Select Parity Group panel (see Figure 5.22) allows you to specify the parity group to be displayed by the **Disk Utility (LDEV)** graph module. The Select Parity Group panel opens automatically after you select the **Disk Utility (LDEV)** module on the Select Module panel.

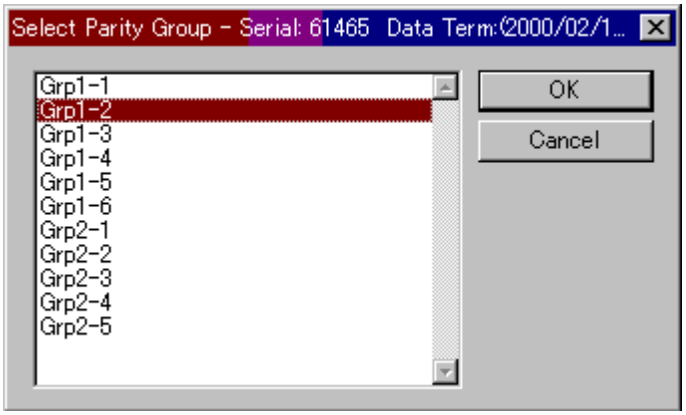

Figure 5.22 Select Parity Group Panel

The Select Parity Group panel displays the available parity groups and allows you to select the desired parity group. The **OK** button applies your selection and opens the Select Term panel (see section 5.5.3). The **Cancel** button cancels your request and returns you to the Select Module panel.

5.5.3 Selecting the Data Term

The Select Term panel (see Figure 5.23) allows you to select the data term, step (data interval), and plotting mode to be displayed on the graphs. The Select Term panel opens automatically when you select **OK** on the Select Module panel (see section 5.5.1) or Select Parity Group panel (see section 5.5.2). The Select Term panel also allows you to change the data term, step, and/or plotting mode and replot a currently displayed graph (see section 5.6.2).

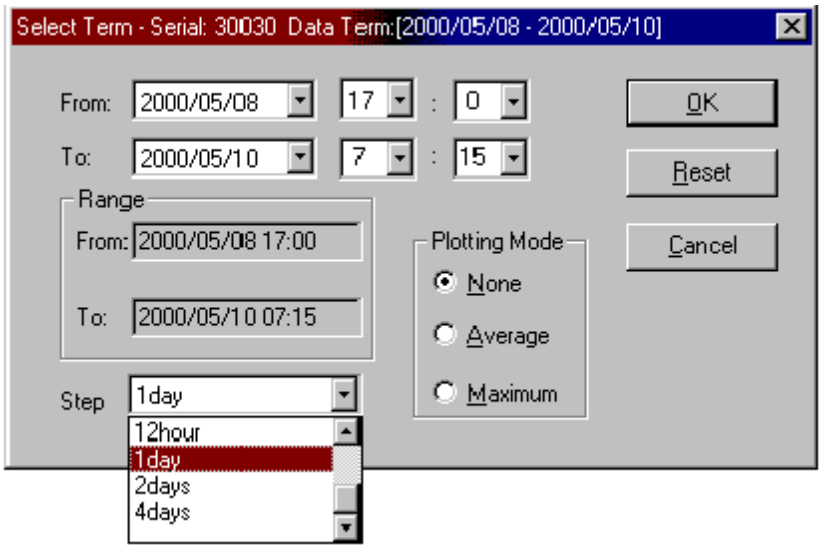

Figure 5.23 Select Term Panel

The **From:** and **To:** combo boxes allow you to specify the term of data to be displayed on the graphs (within the range displayed in the **Range** box). Select the start and end date (year, month, day), hour (0 - 24), and minute (0, 15, 30, 45). ***Note:*** The start time (**From:**) must be earlier than the end time (**To:**). You cannot enter values (e.g., 35) in the minute box.

The **Range** box displays the available range of data for the selected module. The **Step** combo box allows you to specify the scale of the graphs to be displayed: 15 minutes, 30 minutes, 1 hour, 2 hours, 6 hours, 12 hours, 1 day (default), 2 days, or 4 days. The **Plotting Mode** box allows you to specify the graph type as follows:

- **None:** Allows you to plot a graph of the actual usage rate on the scale specified in the **Step** box. For example, if you specify 1 hour, the usage rate collected at times such as 01:00, 02:00, and 03:00 will be plotted.
- **Average:** Allows you to plot a graph of the average usage rate during each interval on the scale specified in the **Step** box. For example, if you specify 1 hour, the average of the data collected during each hour (4 data points per hour) will be plotted.
- **Maximum:** Allows you to plot a graph of the maximum usage rate during each interval on the scale specified in the **Step** box. For example, if you specify 1 hour, the highest data value collected during each hour will be plotted.

The **OK** button loads the statistical data according to the settings on the Select Term panel and displays the Select Unit panel (see section 5.5.4). The **Reset** button resets the Select Term settings to default. The **Cancel** button cancels your request and closes the Select Term panel.

5.5.4 Selecting the Units (LDEVs)

The Select Unit panel (see Figure 5.24) allows you to select the logical devices (LDEVs) to be displayed on the graphs. The Select Unit panel opens automatically when you select **OK** on the Select Term panel (refer to Figure 5.23). The Select Unit panel appears only when the **Disk Utility (LDEV)** module has been selected.

The Select Unit panel also allows you to change the LDEVs and replot the currently displayed graph. To open the Select Unit panel while viewing a data graph, select **Unit...** on the **Configuration** menu, or select the icon on the GraphTool Main panel.

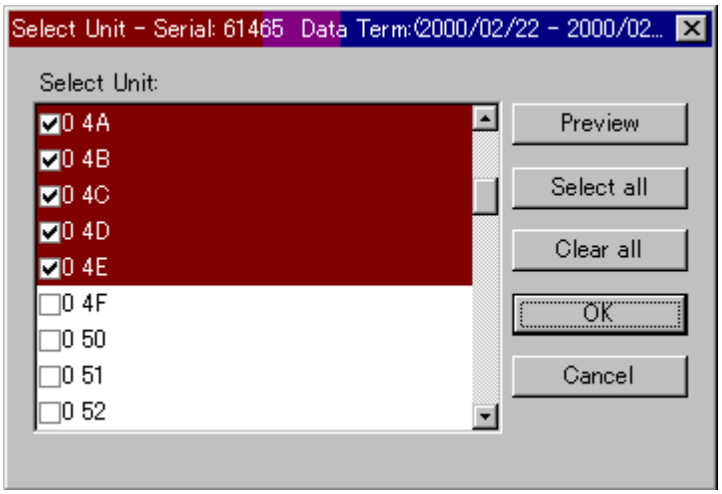

Figure 5.24 Select Unit Panel

The **Select Unit** list box displays the available LDEVs (CU image, LDEV ID). The **Preview** button displays graphs of the selected LDEVs, so that you can check the data and appearance of the graphs in advance. The **Select all** button selects all LDEVs displayed in the **Select Unit** list box. The **Clear all** button deselects all LDEVs displayed in the **Select Unit** list box.

The **OK** button closes the Select Unit panel and displays the **Disk Utility (LDEV)** graphs of the selected LDEVs. The **Cancel** button cancels your request and closes the Select Unit panel.

5.6 Viewing the Graphs

After you have extracted the log data (see section 5.4) and selected the data module and graph parameters (see section 5.5), GraphTool loads and displays the requested data graph. Figure 5.25, Figure 5.26, and Figure 5.27 show examples of the GraphTool data graphs. While data graphs are being displayed, GraphTool allows you to:

- Open additional graphs from the log data file which is currently open (refer to section 5.5),
- Highlight specific data on a data graph (see section 5.6.1),
- Change the data term, step (data interval), and plotting mode of a graph (see section 5.6.2),
- Change the LDEVs being displayed (Disk Utility (LDEV) module only) (see section 5.6.3), and
- Change the font being displayed on the graphs (see section 5.6.4).

Figure 5.25 shows an example of a data graph for the Disk Utility (Parity Group) module. This type of graph is a line graph.

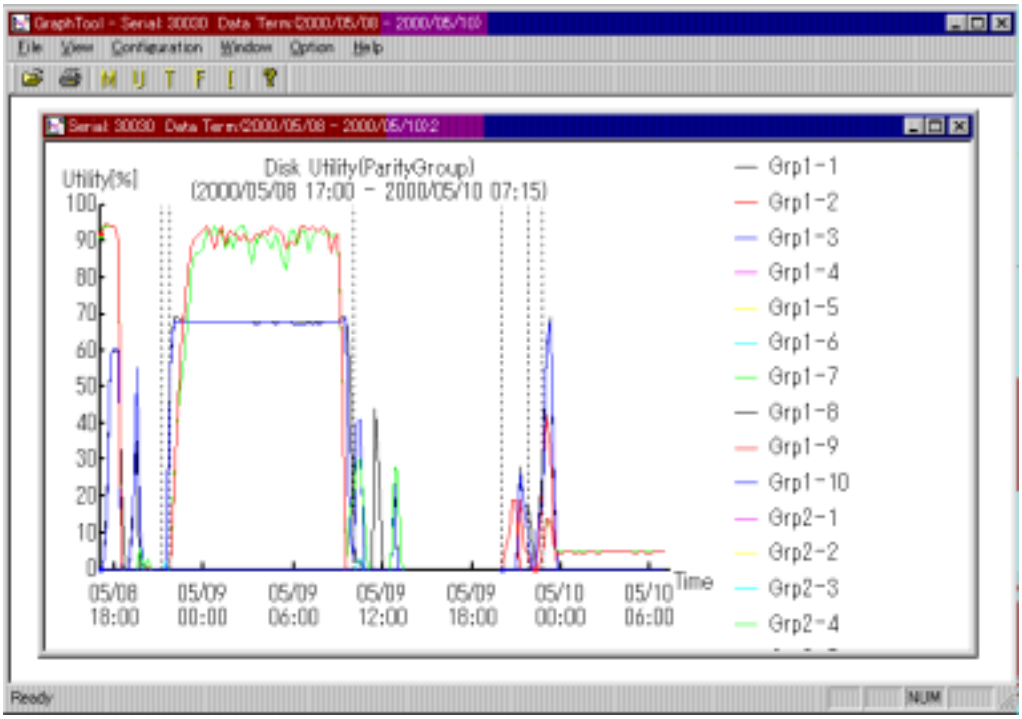

Figure 5.25 Sample Data Graph for Disk Utility (Parity Group) Module

- The vertical axis displays usage rate (utility) from 0% to 100% in units of 10%.
- The horizontal axis displays the date and time of each data sample. The interval between data samples is set according to the **step** selected on the Select Term panel.
- The module (selected on the Select Module panel) is displayed at the top of the data graph.
- The legend is displayed along the right side of the data graph. You can select an item in the legend to highlight the corresponding graph on the left (see section 5.6.1).
- If an LDEV has been relocated, vertical dotted lines show the LDEV's relocation history.

  Note: The data graphs for the 7700E subsystem do not show vertical dotted lines, because the 7700E HIHSM data does not include the LDEV relocation history.

Figure 5.26 shows an example of a data graph for the Disk Utility (LDEV) module. This graph is a stacked area and line graph. An area plotted for each LDEV usage rate, and a line is plotted for the total usage rate of the parity group.

Note: If you specify **Maximum** for the plotting mode on the Select Term panel, the total value may exceed 100%.

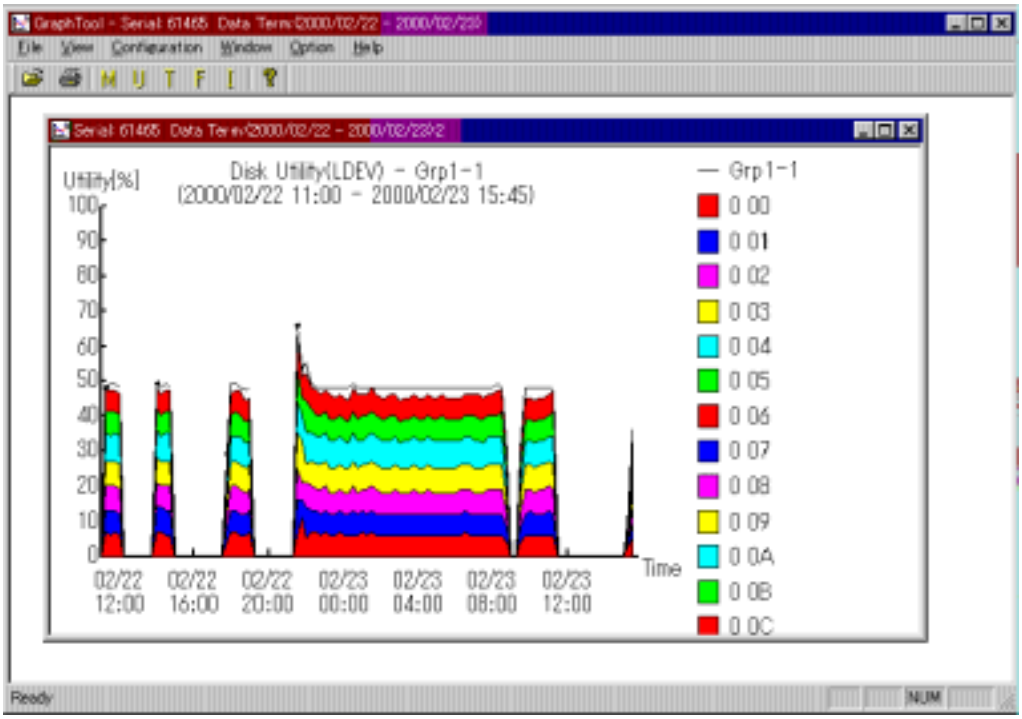

Figure 5.26 Sample Graph for Disk Utility (LDEV) Module

- The vertical axis displays usage rate (utility) from 0% to 100% in units of 10%.
- The horizontal axis displays the date and time of each data sample. The interval between data samples is set according to the **step** selected on the Select Term panel.
- The module and parity group (selected on the Select Module and Select Parity Group panels) are displayed at the top of the data graph.
- The legend is displayed along the right side of the data graph. You can select an item in the legend to highlight the corresponding graph on the left (see section 5.6.1).

  Note: If you display graphs for a 7700E subsystem whose LDEVs have been relocated, the graphs may differ from the actual data, because the 7700E HIHSM data does not include the LDEV relocation history.

Figure 5.27 shows an example of a data graph for the CHP Utility module. The CHP, DKP, DRR, and Path Utility data graphs have the same appearance. This graph is a line graph.

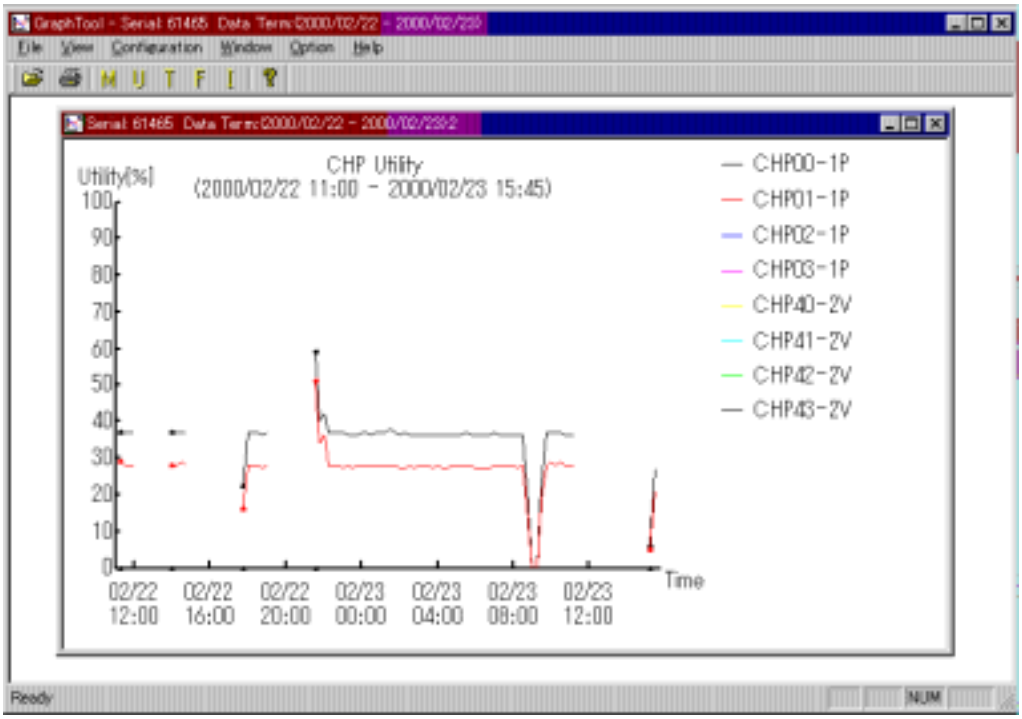

Figure 5.27 Sample Graph for CHP, DKP, DRR, and Path Utility Modules

- The vertical axis displays usage rate (utility) from 0% to 100% in units of 10%.
- The horizontal axis displays the date and time of each data sample. The interval between data samples is set according to the **step** selected on the Select Term panel.
- The module (selected on the Select Module panel) is displayed at the top of the data graph.
- The legend is displayed along the right side of the data graph. You can select an item in the legend to highlight the corresponding graph on the left (see section 5.6.1).

5.6.1 Highlighting Specific Data on a Graph

The legend on each data graph identifies the data being plotted by color. GraphTool allows you to select an item in the legend to highlight the corresponding data on the graph. When you select a legend item, the background color of the item changes to pink. For a line graph, the highlighted line becomes thicker and moves to the foreground of the window (see Figure 5.28). For a stacked area graph, the highlighted stacked area is displayed with a grid pattern (see Figure 5.29). GraphTool also allows you to print graphs showing highlighted data.

Note: The highlight is released when you select the same legend again, or when you select another part of the window, or when you change the LDEVs using the Select Unit panel.

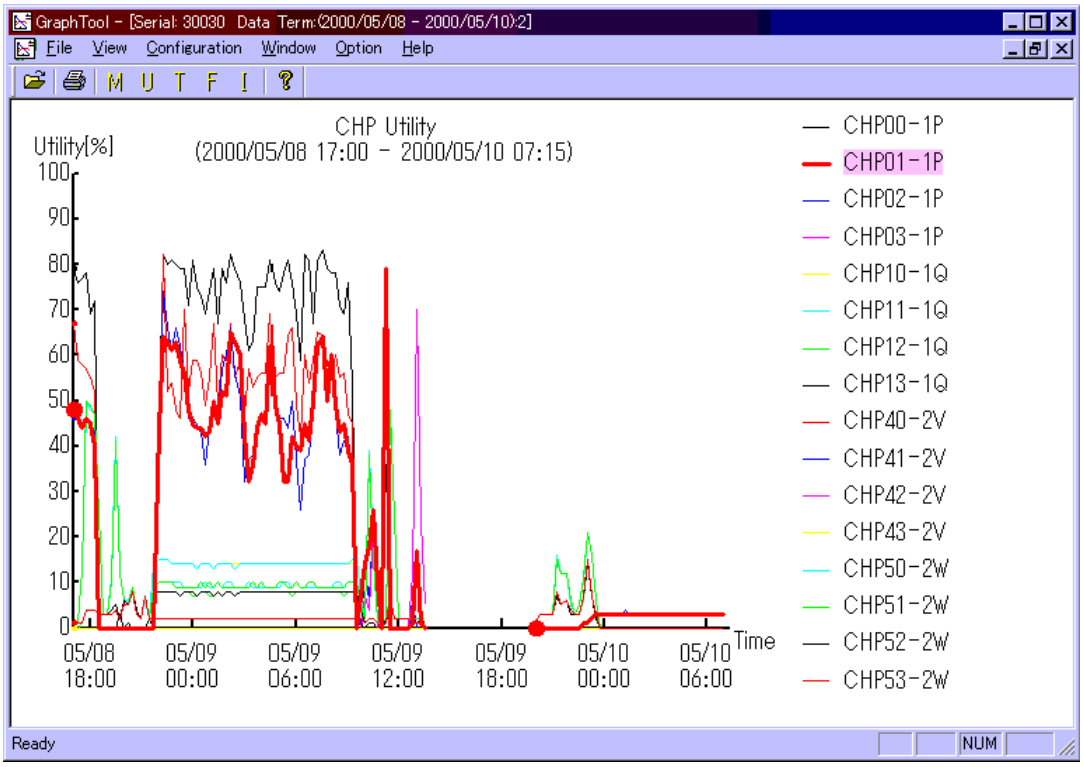

Figure 5.28 Example of a Highlighted Line Graph

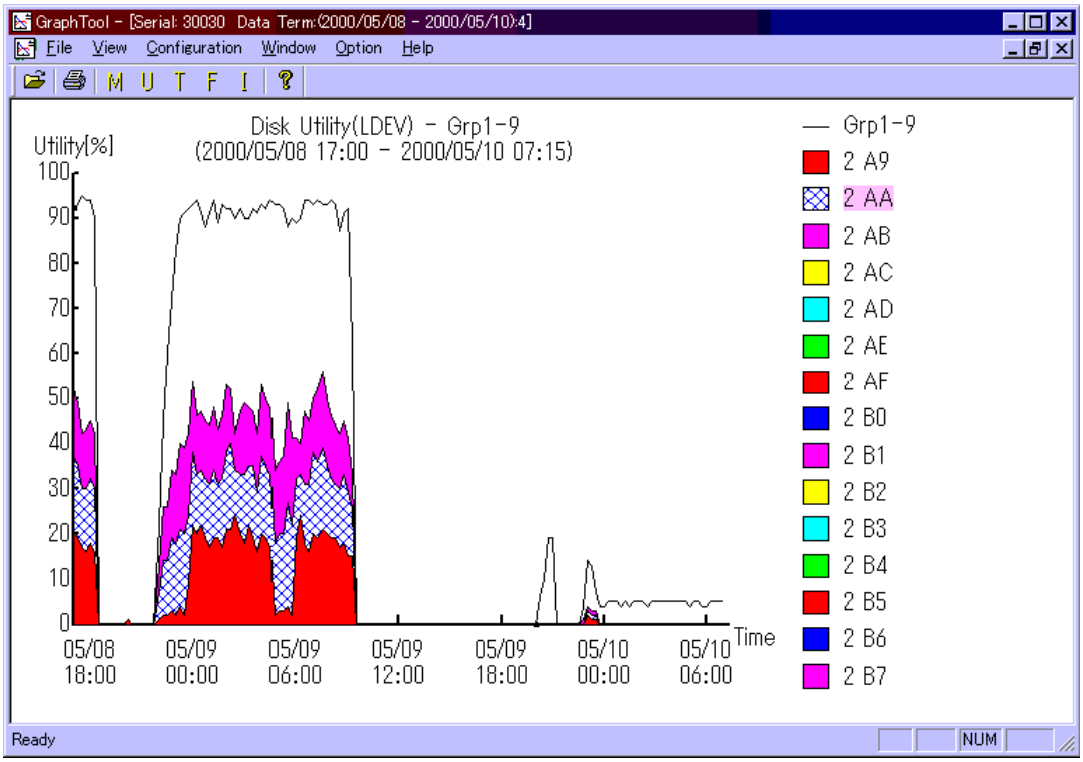

Figure 5.29 Example of a Highlighted Stacked Area Graph

5.6.2 Changing the Data Term, Step, and Plotting Mode of a Graph

The Select Term panel (see Figure 5.30) allows you to change the data term, step (data interval), and/or plotting mode of and replot a currently displayed graph. See section 5.5.3 for a description of the Select Term panel.

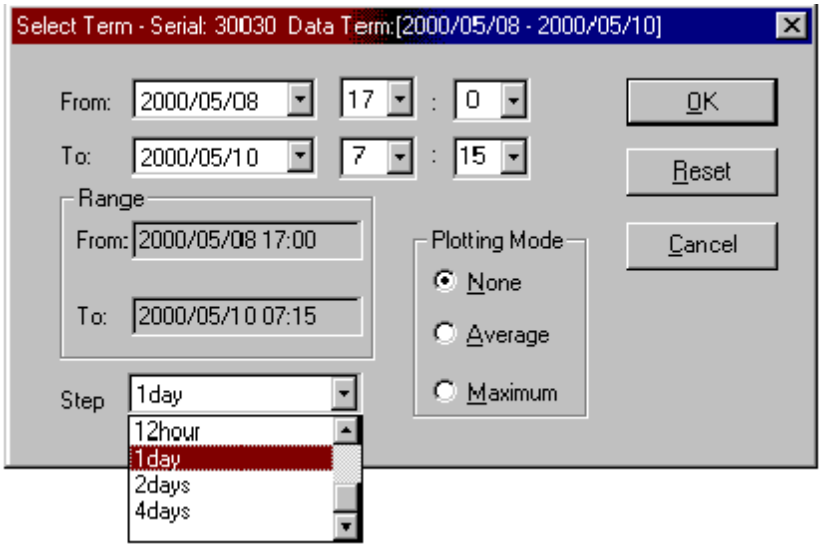

Figure 5.30 Changing the Term, Step, and/or Plotting Mode of a Graph

To change the term parameters for and replot a graph:

1. If more than one graph window is open, make sure that the desired window is active.
2. Open the Select Term panel by selecting **Term** from the **Configuration** menu or by selecting the T icon on the GraphTool Main panel.
3. Change the parameters on the Select Term panel as desired. To reset the parameters to the default settings, select the **Reset** button. To cancel your request and close the panel, select the **Cancel** button.
4. When the desired term parameters are selected, select the **OK** button to close the Select Term panel. GraphTool replots the graph according to the new term settings.

5.6.3 Changing the LDEVs on a Graph

The Select Unit panel (see Figure 5.31) allows you to change the LDEVs and replot a currently displayed graph. See section 5.5.4 for a description of the Select Unit panel.

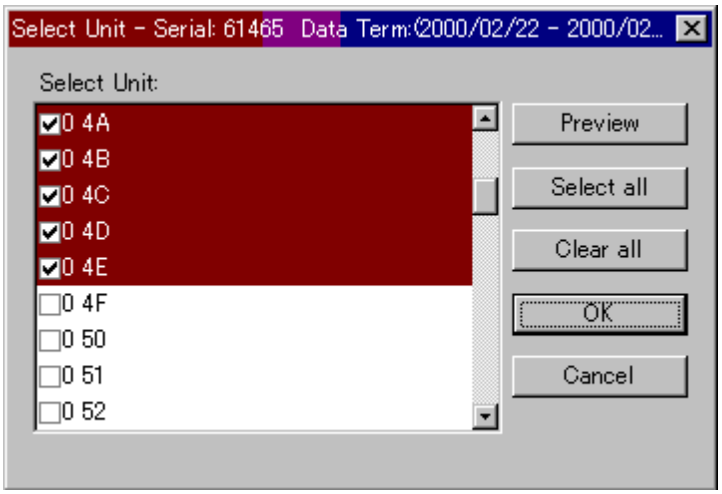

Figure 5.31 Changing the LDEVs on a Graph

To change the LDEVs and replot a graph:

1. If more than one graph window is open, make sure that the desired window is active.
2. Open the Select Unit panel by selecting **Unit** from the **Configuration** menu or by selecting the icon on the GraphTool Main panel.
3. Select the desired LDEVs in the **Select Unit** list box. To select all LDEVs, select the **Select all** button. To deselect all LDEVs, select the **Clear all** button. To check the appearance of the graph before closing the panel, select the **Preview** button. To cancel your request and close the panel, select the **Cancel** button.
4. When the desired LDEVs are selected, select the **OK** button to close the Select Unit panel. GraphTool replots the graph according to the new LDEV settings.

5.6.4 Changing the Font on the Graphs

The Font panel (see Figure 5.32) allows you to change the font properties for and replot a currently displayed graph.

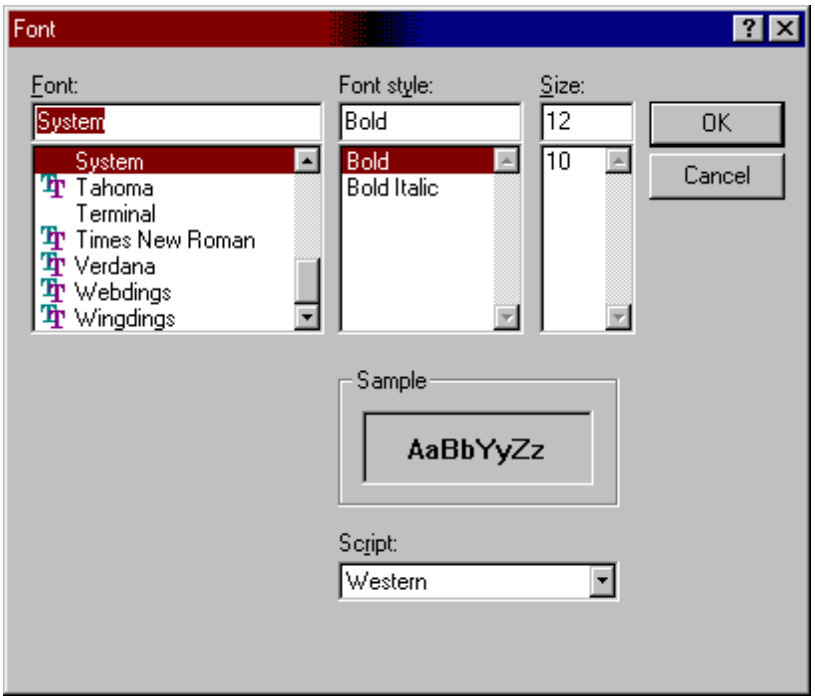

Figure 5.32 Font Panel

To change the font properties for and replot a graph:

1. If more than one graph window is open, make sure that the desired window is active.
2. Open the Font panel by selecting **Font** from the **Configuration** menu or by selecting the F icon on the GraphTool Main panel.
3. Select the desired font type, style, size, and script.To cancel your request and close the panel, select the **Cancel** button.
4. When the desired font properties are selected, select the **OK** button to close the Font panel. GraphTool replots the graph according to the new font properties.

5.7 Printing the Graphs

GraphTool allows you to print the data graphs. You can also change the printer settings by selecting **Print Setup** from the **File** menu on the GraphTool Main panel.

To preview and print a data graph:

1. If more than one graph window is open, make sure that the desired window is active.
2. If you want to display a preview of the print image, select **Print Preview…** from the **File** menu on the GraphTool Main panel. The Print Preview window (see Figure 5.33) displays one or two pages at a time and allows you to zoom in and zoom out. Select **Print…** to print the graph, or select **Close** to close the Print Preview window without printing.
3. If you closed the Print Preview panel without printing, select **Print** from the **File** menu or select the  icon on the GraphTool Main panel to print the graph in the active window.

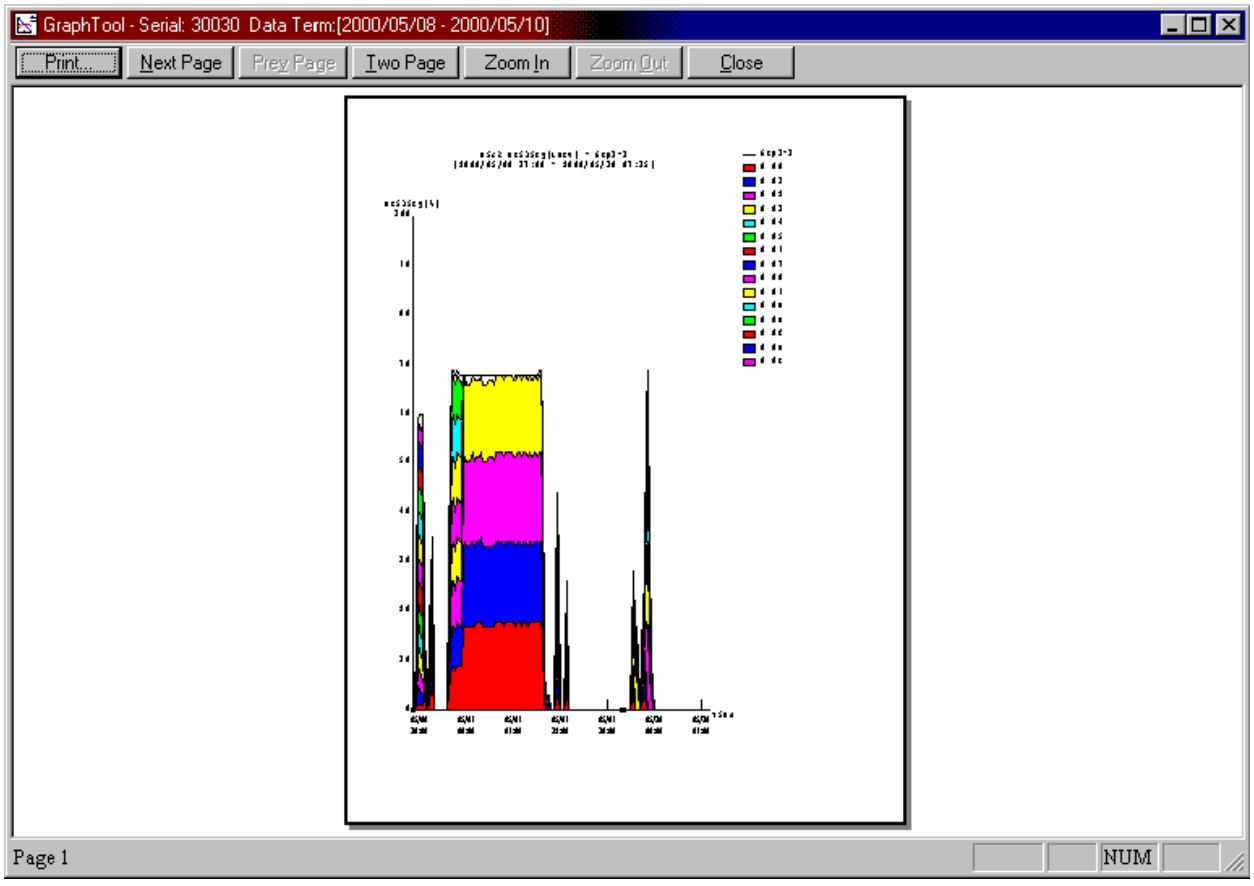

Figure 5.33 Print Preview Window

5.8 Closing the Graphs

To close a data graph window, select the close button (icon at the top left of the window). When you close the last window for the log data file which is currently open, GraphTool will also close the log data file. GraphTool displays a confirmation message for this operation (see Figure 5.34). Select **Yes** to close the last window and also close the log data file. Select **No** to cancel your request to close the last window.

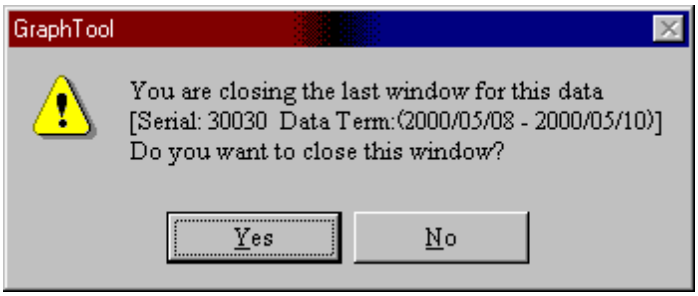

Figure 5.34 Confirmation Message for Closing the Last Window

Chapter 6 Troubleshooting

6.1 Troubleshooting

The Hitachi Lightning 9900™ subsystem provides continuous data availability. For troubleshooting information on the 9900 subsystem, please refer to the *Hitachi Lightning 9900™ User and Reference Guide*, or contact your Hitachi Data Systems representative.

The HIHSM software displays error messages on the Remote Console PC when error conditions occur during HIHSM operations. The error message describes the error and provides a four-digit error code. The first two digits of the error code indicate the error type, and the last two digits provide more specific information about the error. The error message may also include an SVP error code. If you need to call Hitachi Data Systems for assistance, please report the HIHSM and SVP error codes. Please refer to the *9900 Remote Console User's Guide* for a list of error codes displayed on the Remote Console PC.

Table 6.1 (on the next page) lists the HIHSM error codes displayed on the 9900 Remote Console PC and provides recommended actions for resolving the error conditions.

6.2 Calling the Support Center

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including the circumstances surrounding the error or failure and the exact content of any error messages displayed on the host system(s). Please check the remote service information messages (R-SIMs) logged on the Remote Console PC, and note the reference codes and severity levels of the recent R-SIMs. Also note the HIHSM (or other) error code(s) displayed on the Remote Console PC.

The worldwide Hitachi Data Systems Support Centers are:

- Hitachi Data Systems North America/Latin America
  San Diego, California, USA
  1-800-348-4357
- Hitachi Data Systems Europe
  Contact Hitachi Data Systems Local Support
- Hitachi Data Systems Asia Pacific
  North Ryde, Australia
  011-61-2-9325-3300

Table 6.1 HIHSM Error Codes

Code	Description	Recommended Action
0401	A locking time-out was detected during an internal processing. (A retry may result in a normal termination.)	Retry about five seconds later.
0801	The HIHSM could not be used.	Install the PP or the HIHSM.
0810	The issued command cannot accepted in this status. (The command was rejected.)	Check the pair status and confirm whether the command is allowed to be issued.
0811	The issued command is treated as an NOP.	The issued command is treated as an NOP.
0830	A pair cannot be created because the track format is different.	Make sure that the emulation type of the source volume and target volume is the same.
0831	A pair cannot be created because a number of slots is different.	Make sure that the capacity of the source volume and target volume is the same.
0834	The emulation type of the specified source volume is not supported by the HIHSM.	Make sure that the emulation type of the source volume is supported by HIHSM.
0835	The emulation type of the specified target volume is not supported by the HIHSM	Make sure that the emulation type of the target volume is supported by HIHSM.
0836	The pair cannot be created.	Confirms whether the emulation type of the source and the target volume are the same.
0C70	An the source volume is not installed.	Volumes not installed are not included in the object to be processed.
0C71	The source volume cannot be used.	Call the service personnel to make the source volume status normal.
0C72	The source volume is being formatted.	Wait until the formatting of the source volume is completes.
0C73	The source volume is command device.	The command device does not execute HIHSM function
0C80	The target volume is not installed.	Volumes not installed are not included in the object to be processed.
0C81	The target volume cannot be used.	Call the service personnel to make the target volume status normal.
0C82	The target volume is being formatted.	Wait until the formatting of the target volume completes.
0C83	The target volume is command device.	The command device does not execute HIHSM function.
0C90	The LDEV specified to be a Reserve volume is not installed.	Volumes not installed are not included in the object to be processed.
0C91	The LDEV specified to be a Reserve volume cannot be used.	Call the service personnel to make the volume specified to be a Reserve volume normal.
0C92	The LDEV specified to be a Reserve volume is being formatted.	Wait until the formatting of the volume specified to be a Reserve volume completes.
0C93	The LDEV specified to be a Reserve volume is command device.	The command device does not execute HIHSM function.
1009	The multiplicity of the HIHSM exceeded its limit.	Delete some of the pairs.
1011	The number of the LDEV specified to be a Reserve volume has already been used for a Reserve volume.	Change the LDEV number for specifying a Reserve volume.

Table 6.1 HIHSM Error Codes (continued)

Code	Description	Recommended Action
1012	The number of the LDEV specified to be a Reserve volume has already been used for a primary LDEV of the HMRCF/HOMRCF.	Change the LDEV number for specifying a Reserve volume.
1013	The number of the LDEV specified to be a Reserve volume has already been used for a source volume of the hierarchical control.	Change the LDEV number for specifying a Reserve volume.
1014	The number of the LDEV specified to be a Reserve volume has already been used for a target volume of the hierarchical control.	Change the LDEV number for specifying a Reserve volume.
1015	The LDEV specified to be a Reserve volume is not set as a Reserve volume.	Check the volume status.
1017	A Reserve volume cannot be set because the number of Reserve volumes allocated for them was exceeded.	Delete any of the Reserve volumes
102B	The specified volume as the “Reserved” has been set as the HMRCF/HOMRCF target volume.	Check the volume status.
102C	The specified source volume has been set as the HRC/HORC source volume.	Check the volume status.
102D	The specified target volume has been set as the HRC/HORC target volume.	Check the volume status.
102F	The specified source volume has been set as the HIHSM source volume.	Check the volume status.
1030	The specified source volume number does not exist.	Retry after refreshing the screen.
1031	The specified source volume has been set as a Reserve volume.	Check the pair status.
1034	The specified source volume has been set as other source volume.	Check the pair status.
1036	The specified source volume has been set as the HMRCF/HOMRCF source volume.	Delete the pair of the HMRCF/HOMRCF.
1037	The specified source volume has been set as a target volume of the HMRCF/HOMRCF.	Delete the pair of the HMRCF/HOMRCF.
1038	The specified source volume has been set as a M-VOL of the HODM.	Delete the HODM pair.
103B	It is impossible to make a pair because Volume which was specified as a target volume is Root Volume already	It confirms a pair condition.
103C	It is impossible to make a pair because Volume which was specified as a target volume is Node Volume already	It confirms a pair condition.
1040	The specified target volume does not exist.	Execute "Refresh" and do again.
1041	The specified target volume does not set as the Reserve volume.	Check the volume status.
1042	The specified target volume has been used as other HIHSM target volume.	Check the volume status.
1043	The specified target volume has been used as HMRCF/HOMRCF target volume.	Delete pair of HMRCF/HOMRCF
1046	The specified secondary LDEV is used as a P-VOL of the HORC.	Delete the HORC pair.

Table 6.1 HIHSM Error Codes (continued)

Code	Description	Recommended Action
1047	The specified secondary LDEV is used as a P-VOL of the HORC.	Delete the HORC pair.
1048	The specified secondary LDEV is used as an M-VOL of the HODM.	Delete the HORC pair.
104A	The specified secondary LDEV is used as a source volume of the HMRCF / HOMRCF.	Delete the HMRCF/HOMRCF pair.
104B	The specified target volume has been used as other HIHSM source volume.	Check the volume status.
104C	The LDEV specified as a Reserve volume is being used as a HODM M-VOL.	Delete the HODM pair.
104E	The LDEV specified as a Reserve volume is being used as an M-VOL of the HRC/HORC.	Delete the HRC/HORC pair.
104F	The LDEV specified as a Reserve volume is being used as an R-VOL of the HRC/HORC.	Delete the HRC/HORC pair.
1051	The LDEV numbers of the specified source and the target volumes are the same.	Retry after refreshing the screen.
1085	The source volume was used for Reserve volume of HIHSM.	Check the source volume or Cancel the Reserve volume of HIHSM.
1088	The specified reserve volume has been set as DCR.	Reset the DCR.
1089	The specified source volume has been set as DCR.	Reset the DCR.
108A	The specified target volume has been set as DCR.	Reset the DCR.
108B	The specified reserve volume has been constructed as LUSE.	Reset the LUSE.
108C	The specified reserve volume has been set as HMRCF/HOMRCF reserved volume.	Check the pair status.

Appendix A Acronyms and Abbreviations

ACP	array control processor
ave	average
CARB	cache memory arbiter
CCI	Hitachi Command Control Interface
CHA	channel adapter
CHIP	client-host interface processor
CHP	channel processor
CSW	cache path switch
CU	control unit
CVS	Custom Volume Size (also called Virtual LVI and Virtual LUN)
DCR	dynamic cache residency
DKA	disk adapter
DKC	disk controller
DKP	disk processor
DLL	dynamic link library
DRR	data recovery and regeneration processor
dst	destination (target)
grp	group
HDD	hard disk drive
HIHSM	Hitachi Internal Hierarchical Storage Manager
HMRCF	Hitachi Multi-RAID Coupling Feature (ShadowImage for mainframe)
HODM	Hitachi Online Data Migration
HOMRCF	Hitachi Open Multi-RAID Coupling Feature (ShadowImage for open systems)
HORC	Hitachi Open Remote Copy
HRC	Hitachi Remote Copy
I/O	input/output
LAN	local-area network
LDEV	logical device
LU	logical unit (e.g., OPEN-3)
LUN	logical unit number
LUSE	Logical Unit Size Expansion
LVI	logical volume image (e.g., 3390-3R)
max	maximum
min	minimum, minutes
num	number
PC	personal computer system
R-SIM	remote service information message
RAID	redundant array of independent disks

RCU	remote control unit
RMCMAIN	remote console main
rnd	random
sec	seconds
seq	sequential
SIM	service information message
SM	shared memory
src	source
SVP	service processor
TID	target ID
vol	volume
VOLSER	volume serial number

Index