Intelligent Computer Peripherals ®
ICP

Chapter K

The Diagnosis Program GDTMON

K. The Diagnosis Program GDTMON

GDTMON (GDT monitor) is a helpful and flexible diagnosis tool for the monitoring, maintenance and tuning of mass storage subsystems which are based on one or more ICP Controllers. The key features of GDTMON:

- Diagnosis program with a graphical user interface. Clear performance representation with variable horizontal bars
- Available under MS-DOS, NetWare 3.x & 4.x, Windows 95, Windows NT, OS/2 and SCO UNIX
- Loadable locally (on the server) or remotely from an authorized workstation (NCPE and NETBIOS protocol support)
- Indicates the performance, expressed in [KB/sec] and [IO/sec], of:
  - Host Drives (Disk, Chain, RAID 0, 1, 4, 5, 10)
  - Logical Drives
  - SCSI Drives
- Indicates the GDT cache utilization
  - Read-Hits
  - Write-Hits
  - Separate indication for data- and parity-cache (RAID 4/5)
- Allows online changing of the GDT cache parameters:
  - Cache ON/OFF
  - Delayed Write ON/OFF
- Allows online changing of the SCSI devices' parameters:
  - SCSI protocol
  - Synchronous or Asynchronous data transfer
  - Setting of the synchronous data transfer rate
  - Disconnect/Reconnect
  - Disk Cache ON/OFF
  - Tagged Queues ON/OFF
- Indicates the structure of Logical, Array and Host Drives
- Performs online parity checking of RAID 4 and RAID 5 Host Drives
- Allows the online RAID Level migration and capacity expansion of existing Array Drives
- Saves all relevant configuration data to floppy disk or hard disk
- Sets up or removes RAID 1 Disk Arrays while maintaining full operational conditions
- Performs Hot Plugs on RAID 4 / RAID 5 disk arrays
- Allows to add or remove a dedicated or pool Hot Fix drive while maintaining full operational conditions

K.1 Loading GDTMON

As mentioned before, the GDTMON program is available for various operating systems. It can be used either locally or remotely. This means that all ICP Controllers in a network can be monitored and serviced from one (or several) workstation(s). The communication between the ICP Controller(s) and the GDTMON program is based on the NETBIOS or NCPE protocols. Thus, for example, it is possible to monitor the ICP Controller in an OS/2 workstation or the ICP Controller in the Novell NetWare fileserver from a Windows 95 workstation.
In order to be able to access a certain computer system with its ICP Controller remotely on the network, a special communication module has to be loaded first, which routes the information through the network. For a Novell NetWare fileserver this is the CTRLSRV.NLM module (NCPE). For Windows NT (server and workstation), Windows 95 and OS/2 (server and workstation) this is the NBSERVx.EXE module (NETBIOS). While the module is loaded it searches for a CTRLSRV.CFG (Novell NetWare) or NBSERV.CFG (Windows NT, Windows 95 and OS/2) file, which includes the definition of the access rights of the different users and their passwords. The CFG file assigns every user to two different access levels. Access level 0 gives the user all functions to view and change the controller-, disk-drive and disk-array-settings. Access level 1 entitles the user only to view the various settings and performance data.
The Windows NT, Windows 95 and OS/2 driver diskettes include so-called DLLs for the supported protocols. Example: If the MON4NETB.DLL file is located in the same directory as GDTMON.EXE file for Windows NT, the NETBIOS protocol is automatically used for the GDTMON on this system. As soon as the NETBIOS support of this operating system is installed during the network configuration, the GDTMON on this system can communicate through this protocol with another system in the network which has a ICP Controller.
The following drivers for the 'remote' GDTMONitor are currently available:

OS/2

GDTMON32.EXE	GDTMONitor for OS/2
MON2NETB.DLL	NETBIOS DLL for OS/2
MON2NCPE.DLL	NCPE DLL for OS/2
NBSRV2.EXE	NETBIOS Server for OS/2
NBSRV.CFG	Configuration file for NetBios Server
NBCLEAN2.EXE	NETBIOS analysis program

Windows NT

GDTMON.EXE	GDTMONitor for Windows NT
MON4NETB.DLL	NETBIOS DLL for Windows NT
NBSRV4.EXE	NETBIOS Server for Windows NT
NBSRV.CFG	Configuration file for NETBIOS Server
NBCLEAN4.EXE	NETBIOS analysis program

Windows 95

GDTMON.EXE	GDTMONitor for Windows 95
MON5NETB.DLL	NETBIOS DLL for Windows 95
NBSRV5.EXE	NETBIOS Server for Windows 95
NBSRV.CFG	Configuration file for NETBIOS Server
NBCLEAN5.EXE	NETBIOS analysis program

The MS-DOS GDTMON also allows remote access. The DLLs are integrated into the GDTMON.EXE file, thus there is no NETBIOS server available for MS-DOS. Since there is no NETBIOS support within MS-DOS, it is necessary to load the NETBIOS program to access from a MS-DOS workstation through the network a ICP Controller in a Windows NT and

OS/2 workstation or server, or a Windows 95 workstation. NETBIOS is part of the Novell NetWare operating system. For the remote access of a Novell NetWare fileserver the NETBIOS program is not needed.

K.1.1 Loading the GDTMON Program Under NetWare

The GDTMON program for NetWare is part of the GDT *Novell NetWare* disk.
GDTMON can be used either under NetWare 3.1x or under NetWare 4.0x. There are two different methods of loading GDTMON:

- loading GDTMON on the fileserver
- loading GDTMON on an authorized workstation (remote)

Loading GDTMON on the fileserver. Beforehand, the GDT NetWare driver (GDTRP311.DSK for NetWare 3.11, GDTRP312.DSK for NetWare 3.12 and GDTRP400.DSK for NetWare 4.x) and the auto-loading module CTRLTRAN.DSK must have been loaded on the fileserver.

```
LOAD  GDTMON <ENTER>
```

on the fileserver.

Loading GDTMON on a workstation. In this case, too, the GDT NetWare driver and the auto-loading module CTRLTRAN.DSK must have been previously loaded on the fileserver console. In addition, the module CTRLSRV.NLM has to be loaded. This module searches for a file named CTRLSRV.CFG. This file must be located in the same directory as CTRLSRV.NLM. The system administrator has to set up a user group named GDT_OPERATOR. All users belonging to this group are given access (through GDTMON) to the ICP Controller(s) in this specific fileserver (Access level 0). Now, the GDTMON program can be loaded from one (or more) workstation(s):

```
GDTMON <ENTER>
```

By selecting the menu *Select Controller* of the GDTMON main menu, you can now choose either a fileserver (equipped with a ICP Controller), or a ICP Controller in your workstation.

K.1.2 Loading the GDTMON Program Under OS/2

The GDTMON program for OS/2 is part of the GDT *OS/2* disk. To load the program under OS/2, enter:

```
GDTMON32  <ENTER>
```

K.1.3 Loading the GDTMON Program Under Windows NT

The GDTMON program for Windows NT is part of the *GDT Windows NT* disk. To load the program under Windows NT, enter:

```
GDTMON  <ENTER>
```

K.1.4 Loading the GDTMON Program Under Windows 95

The GDTMON program for Windows 95 is part of the *GDT Windows 95* disk. To load the program under Windows 95, enter:

GDTMON <ENTER>

K.1.5 Loading gdtmon under SCO UNIX

In order to be able to use the gdtmon program under SCO UNIX (2.x, 4.x and 5.x), it becomes necessary to substitute the standard terminal entry by a new one:

```
cd /usr/lib/terminfo <ENTER>
tic  gdt386.src <ENTER>
```

Before each loading of gdtmon, this terminal has to be activated by:

```
TERM  =  gdt386 <ENTER>
export   TERM <ENTER>
```

These two lines can also be inserted in the ***.profile*** file and will then be automatically processed during each login. The gdtmon program itself is copied during the SCO UNIX installation (chapter G of this User's Manual) into the */etc* directory. gdtmon is loaded by entering:

```
gdtmon <ENTER>
```

K.2 The GDTMON Program

As mentioned before, the GDTMON program appears identical for all operating systems, so we can demonstrate the use and functioning of this program regardless of the operating system used. In previous chapters we have already described the hierarchical structure of the GDT firmware. We have defined 4 different levels of hierarchy: Level 1 where the physical SCSI drives named ***Physical Drives*** are found, level 2 containing the ***Logical Drives*** (made up of one or several *Physical Drives*), level 3 where we have the ***Array Drives***, and finally, level 4 where the ***Host Drives*** are. Only the latter ones are known to the operating system. The drive of a given level of hierarchy is always set up by using the drives of the next lower level as components. Accordingly, GDTMON has various menu options, each referring to one level of hierarchy

Host Drives	———	Level 4
Logical Drives	———	Level 2
Physical Drives	———	Level 1

Each menu option displays the performance of the drives belonging to the corresponding level. (Note: The performance of Array Drives and Host Drives is identical). The performance is measured in KB/s (kilobyte per second, transfer rate) and IO/s (I/Os per second, number of simultaneously processed I/Os on the ICP Controller). The performance is displayed numerically as well as graphically in the form of variable horizontal bars, with separate indications for each drive and its write and read accesses.

K.2.1 Select Controller

This menu option yields a list of available ICP Controllers. By selecting a protocol, you have either access to the ICP Controller(s) in your local computer (e.g., MS-DOS, Windows NT, Windows 95 or OS/2 protocol), or to a ICP Controller in another computer (server or workstation) in the network (NetWare NCPE or IBM NETBIOS protocol). All diagnosis and maintenance functions of GDTMON refer to the ICP Controller you have selected here (and the SCSI devices connected with this controller).

(MS-DOS) in the Select Protocol menu indicates that the GDTMON program was loaded on a MS-DOS computer. I.e.: If GDTMON had been loaded under Windows NT, we would see there (Windows NT).

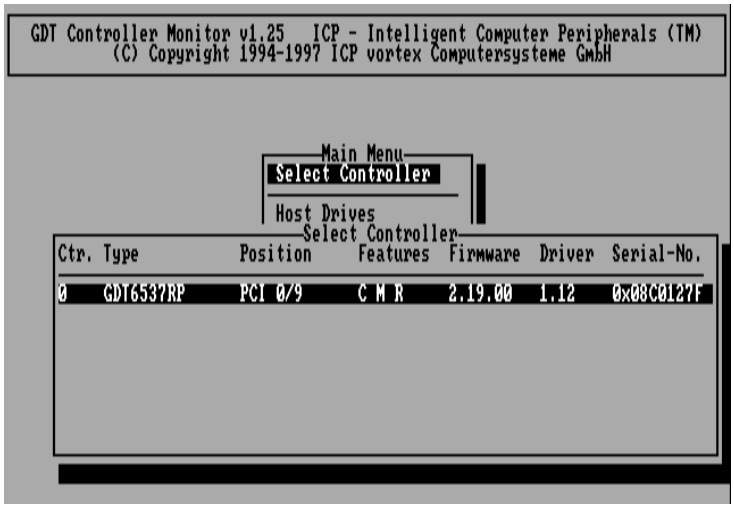

K.2.2 Host Drives

This menu option leads to the list of available Host Drives (level 4). We would like to recall that the operating system (e.g., NetWare) only recognizes these Host Drives and not their possibly complex structures. This means that it is of no importance for the operating system if a Host Drive consists of one single SCSI drive (of the type *disk*), or of 5 SCSI drives configured to form a RAID 4 Array Drive. Apart from the performance, the name, type, state and capacity (1024KB = 1MB) of a Host Drive is displayed. The figures shown at *Total* represent the overall performance of the Host Drives as a whole. With the ← and → keys you may change the scale of the graphical KB/s indication. With the ↑ and ↓ keys you can scroll the screen to see further Host Drives (if available).

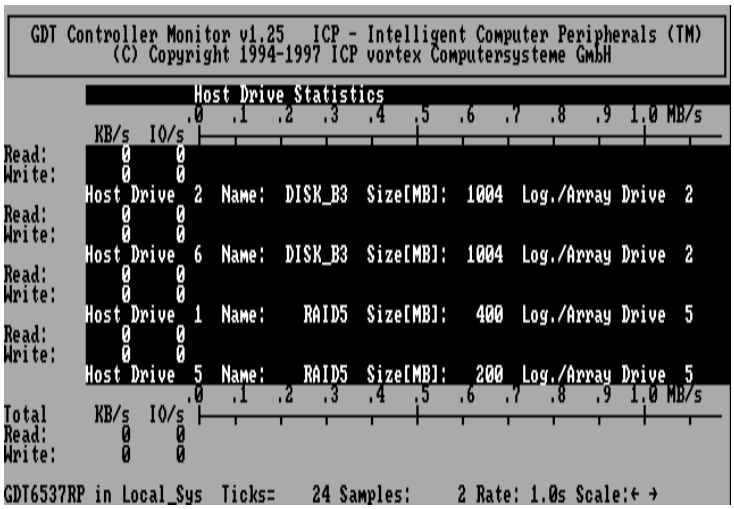

(Note: All Host Drives on the screen shown above are idle.)

K.2.3 Logical Drives

This menu option yields a list of available Logical Drives (level 2). Logical Drives are the components for Array Drives and Host Drives. In its most simple form, a Host Drive consists of one Logical Drive which is made up of a single SCSI drive (type *disk*). In case of RAID Host

Drives, the performance of the Logical Drives forming a RAID Host Drive are shown in the menu *Logical Drives*. The performance data displayed here gives immediate information on the quality of a given Host Drive. When judging the indicated performance of Logical Drives belonging to a RAID Host Drive, the following considerations should be taken into account:

(a) RAID 0, RAID 1, RAID 10 and RAID 5 Host Drives
If a certain Logical Drive shows poor performance for a longer period (when compared to the other Logical Drives), this Logical Drive impairs the overall performance, making it the *bottle neck* of the entire RAID Host Drive.

(b) RAID 4 Host Drives
In principle, the same as in (a) applies to RAID 4 Host Drives, with the exception made for the parity Logical Drive. The poor performance can have various reasons: the SCSI drive forming the Logical Drive

- is too slow (different from the other SCSI drives with regard to type or manufacturer)
- has wrong parameters (SCSI II, disk cache, tagged queues etc.)
- has too many defects, causing the read/write actuator to have to move to *alternate tracks* very often

When setting up RAID 4/5 Host Drives, please observe the notes on RAID 4/5 Host Drives exposed in the chapter **GDTSETUP in Detail**.

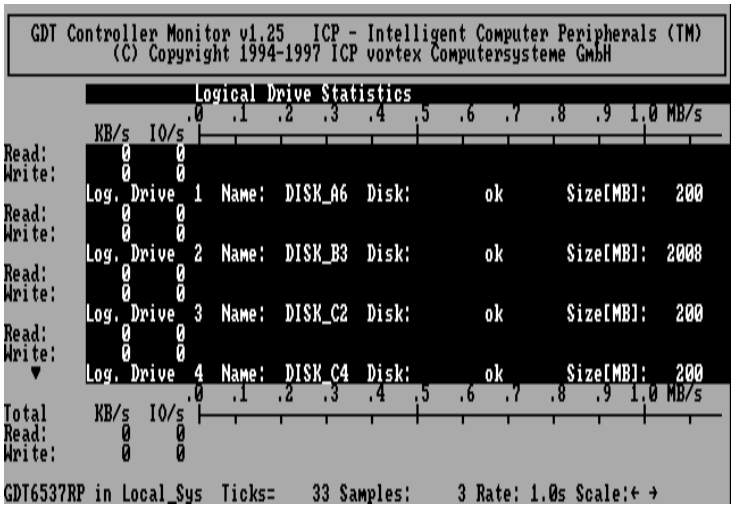

The figures under *Total* represent the total performance of all Logical Drives. With the ← and → keys you may change the scale of the graphical KB/s indication. With the ↑ and ↓ keys you can scroll the screen to see more Logical Drives (if available).

K.2.4 Physical Drives

In addition to the performance report on the SCSI drives, you are given additional information on each device:

- the GDT SCSI channel the SCSI drive is connected to
- which SCSI-ID the drive has
- the name of the SCSI drive
- the gross capacity (1MB = 1024KB)

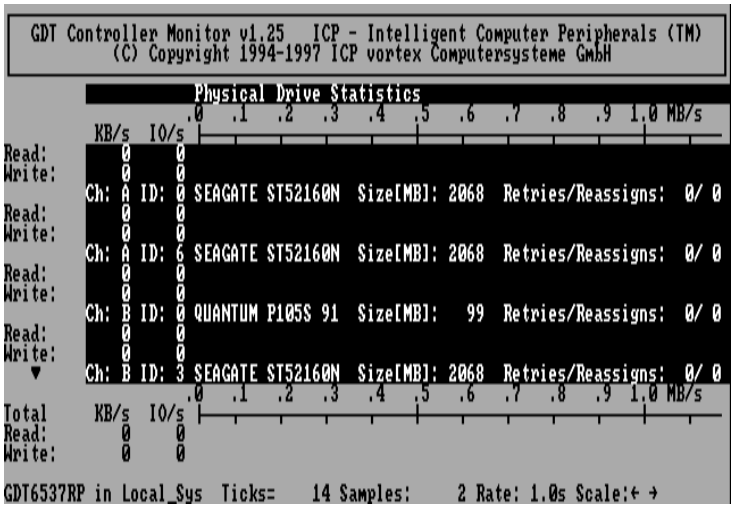

The *Retries/Reassigns* counters have a particular meaning:

(1) The *Retries* counter is incremented by one unit whenever the ICP Controller retries to access a SCSI device. If this counter continues to increase (possibly on other SCSI drives, too) it is very likely that the SCSI cable is not *good* enough for the selected data transfer rate (cable too long, poor quality of cable and connectors), or that the SCSI bus is not properly terminated (too many terminators on the cable, or missing terminator). In very few cases is the SCSI drive concerned defective. The retry counter also increases when the SCSI parameters of a SCSI drive are changed (see further ahead). Obviously, retries due to this do not imply bad SCSI cabling.

(2) The reassign counter reflects the number of media defects which occur on the SCSI drive. Defective blocks of the SCSI device are assigned substitute blocks (spare blocks) which are either on the same track, or on alternate ones if all spare blocks on the same track are already in use. The administration of the reassignments is carried out by the SCSI drive through according reassignment tables. Note: If a SCSI drive works with alternate tracks, it is generally no longer suitable for applications with high performance expectations. Whenever a defective block is being accessed, the read/write actuator has to move to an alternate position and this requires extra time.
If you observe that the number of reassigns is constantly increasing, you may suspect that something is wrong with this drive.
With the ← and → keys, you can change the scale of the graphical KB/s indication.
With the ↑ and ↓ keys, you can scroll the screen to see more SCSI drives (if available).

K.2.5 Cache Statistics

This menu option gives information on the utilization of the GDT cache. For RAID 4/5 Host Drives, the data cache and parity cache are displayed separately.
Note: The GDT firmware only allocates cache RAM to the GDT parity cache if RAID4/5 Host Drives have been set up.

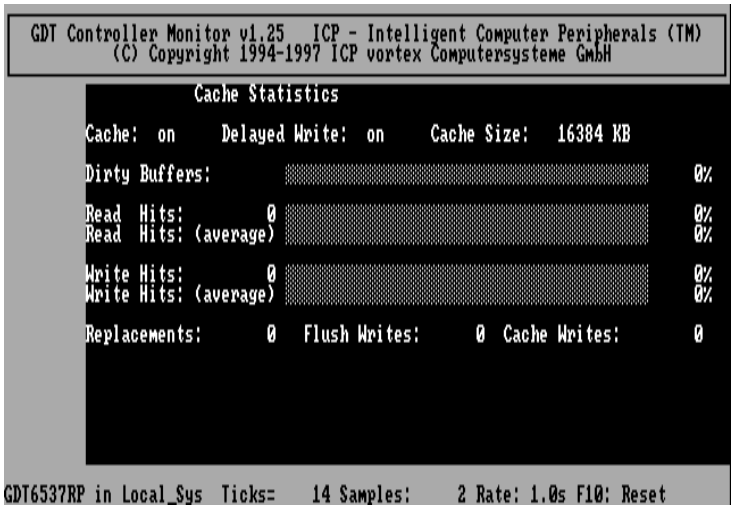

K.2.6 Sampling Rate

By setting the sampling rate, you can choose the interval at which the ICP Controller delivers new measurements. According to the operating system used, the sampling rate can be set to a maximum of 60 seconds. The default setting is 1 second.

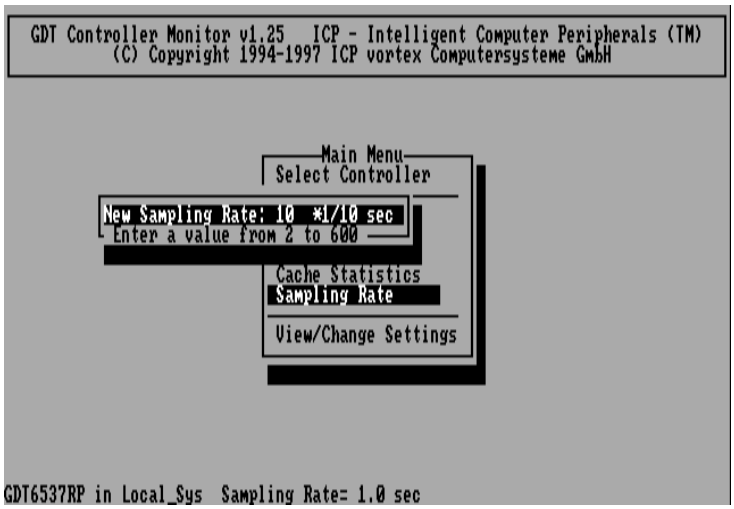

K.3 The Menu: View/Change Settings

This menu includes a set of very powerful options and functions for the online maintenance and diagnosis of RAID 1/4/5/10 Host Drives.

From this menu you can select further submenus:

Controller - View the ICP Controller configuration
- View Last Events
- View/Change the Intelligent Fault Bus settings

Cache Settings - View/Change the GDT cache parameters

Physical Drives - View/Change the SCSI parameters

Logical Drives - Display the structure of Logical Drives
- Add/Remove Mirror Drives to/from Host Drives
- Perform a Hot Plug on a RAID 1 Host Drive
- Add/Remove Private Hot Fix and Pool Hot Fix drives to/from a RAID 1 Host Drive

Array Drives - Display the structure of RAID 4/5/10 Host Drives
- Perform a Parity Verify on RAID 4/5 Host Drives
- Perform a Parity Recalculation on RAID 4/5 Host Drives
- Perfrom online capacity expansion and/or RAID level migration
- Perform a Hot Plug on RAID 4/5/10 Host Drives
- Add/Remove Private Hot Fix and Pool Hot Fix drives to/from RAID 4/5/10 Host Drives

Save Information - Save all relevant configuration data to disk

Before going into detail, some information on the so-called Hot Plug of fault tolerant disk arrays (RAID 1/4/5/10) is given below.

K.3.1 Notes and Information on the Hot Plug function of GDTMON

Fault tolerant means that a hard disk which is part of a RAID 1/4/5/10 Array Drive can fail without causing data loss on the Array Drive. At the same time, the Array Drive remains fully accessible. Obviously, the Array Drive then lacks the redundant data, therefore the defective disk should be replaced by an intact one as soon as possible. In chapters C (*RAIDYNE Quick-Setup*) and L (*GDTSETUP in Detail*), we have shown with various examples how to exchange a defective drive with GDTSETUP. Moreover, we described the functioning of the so-called Hot Fix drive. This "constantly available spare part when needed" automatically integrates itself into the Array Drive and is therefore the quickest means of regaining a redundant Array Drive.
The Hot Plug function enables the replacement of a drive of an Array Drive (either in the *ready* or *fail* state) while the system continues to run, that is, without having to shut down the NetWare fileserver for instance. A drive replacement may not only become necessary when the drive has already failed, but also when there are signs that a failure could occur soon (strong whistling of the hard disk, or constant retries of the read/write head).
Only those users with thorough knowledge of RAID and the ICP Controller should use the Hot Plug function. Improper use can lead to data loss. (Naturally, we have integrated all kinds of security provisions into the ICP Controller and GDTMON. But how can we prevent a user from plugging out the wrong drive?). We recommend that you document each Array Drive immediately after its configuration with GDTSETUP. This record should at least comprise the following information:

- To which ICP Controller has the Array Drive been connected ?
- Which SCSI devices (hard disks) are part of the Array Drive ?
- To which SCSI ID have they been set ?
- Which SCSI devices terminate a SCSI channel ?

In addition, the SCSI devices themselves should be labeled with the above information. The following is an example of such a label.

C xy z t

C: controller number
x: SCSI channel
y: SCSI-ID
z: Logical Drive number
t: **+** = terminated, **-** = not terminated

Example:

1 B6 3 +

controller number 1, SCSI channel B, SCSI-ID 6, Logical Drive number 3, terminated. The menu option *Save Information* of GDTSETUP, GDTMON program can be of help when setting up the documentation.

How does the Hot Plug mechanism work ?

During the Hot Plug, the SCSI channel of the ICP Controller to which the drive to be exchanged is connected, is temporarily "halted" (for the time needed for the exchange), so that the drive can be disconnected from the SCSI bus and replaced without any risk. After the replacement, the SCSI channel halt is lifted and the firmware automatically begins to

rebuild the new drive. The halting and the halt-lifting of the GDT SCSI channel is controlled by the Hot Plug function. The above mentioned halting of the SCSI channel avoids that interferences caused by plugging off the drive impair the functioning of other drives still active on this SCSI channel. Moreover it prevents the possible destruction of the SCSI processors of the ICP Controller or the SCSI protocol ICs of the SCSI drives. However, this implies that none of the SCSI devices of the halted SCSI channel can be accessed during the time the defective drive is being exchanged. This may affect, for example, other drives of this Array Drive, other Array Drives, simple Host Drives (of the type disk etc.), and *Not Direct Access Devices* (CD-ROMs, DAT streamers, MOs, etc.). If the ICP Controller has only one SCSI channel, or if all SCSI devices are connected to one single SCSI channel, then no SCSI device can be accessed during the time of replacement. Therefore, it is evident that it is best if the ICP Controller had as many SCSI channels as possible and that all SCSI devices were distributed equally to the available channels in order to avoid that the Array Drive or other SCSI devices cannot be accessed during the Hot Plug drive replacement. This non-accessibility of the Array Drive during the replacement is communicated to the operating system by the GDT driver program that integrated the ICP Controller into the operating system. If the period of non-accessibility becomes too long a system error may occur. The Hot Plug should be carried out as quick as possible.

Various manufacturers offer so-called *Disk Shuttle* (sometimes also called *Disk Shelf*) subsystems. There are two categories of these subsystems: Non-Intelligent and Intelligent. Both generally consist of the shuttle itself, a shock safe metal or plastic enclosure containing the hard disk, a frame which is mounted in the computer case or in an external subsystem, and a sufficient cooling system.
Especially the last component is very important. If the hard disks run too hot, it is very likely that they will fail. In addition, their lifetime is reduced dramatically. The mounting frame has appropriate slide rails and locks that fit with the shuttle. The shuttle can be slid into or out of the mounting. When using any of these systems you have to make sure that the SCSI termination (regardless of whether the drive has been terminated or not) is made in the mounting frame. In addition, always use high quality connectors for the SCSI buses and the power supply. It is very important that during the plug out, the SCSI bus is disconnected first, and then the power supply. When plugging in again, the order is vice versa.
The Non-Intelligent subsystems have no communication with the ICP Controller and no intelligent backplane (i.e., with sensors).

If it is intended to avoid this halting and halt-lifting of the GDT SCSI channel, an Intelligent subsystem is needed. The ICP Controllers support two different types of such subsystems: SAF-TE and DEC(TM) fault bus compliant subsystems. SAF-TE stands for SCSI Accessed Fault-Tolerant Enclosure and is soon becoming part of the SCSI specification. The DEC(TM) fault bus is a special technology from Digital Equipment (TM).
With these subsystems, the ICP Controller enables the performance of the so-called *Auto Hot Plug*, which reduces the whole Hot Plug procedure to the simple replacement of the disk shuttle which contains the desired hard disk. Furthermore, these subsystems have typically several LEDs to indicate for example the disk shuttle which contains the defective hard disk.

K.3.2 Controller

This option displays details on the ICP Controller. For example, how much Cache RAM the ICP Controller has and what the current termination setting of the various SCSI channels is.

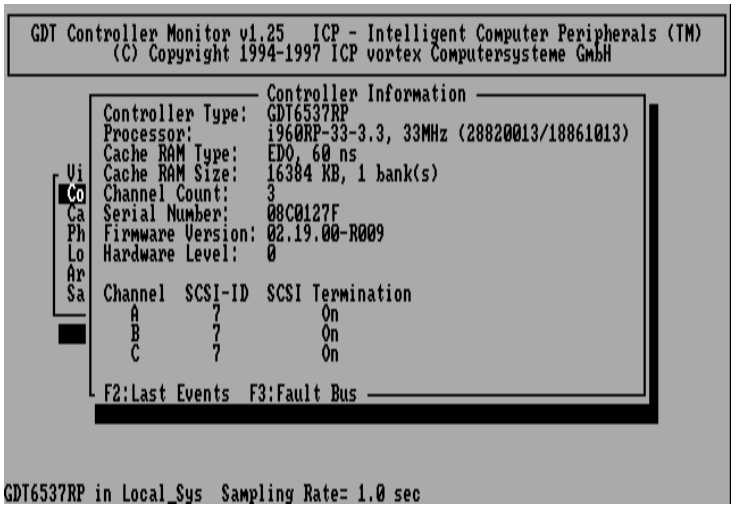

Press <F3> to get detailed information on the configuration of the Intelligent Fault Bus (IFB). Changing and enabling these features requires the existence of an IFB-compatible subsystem.

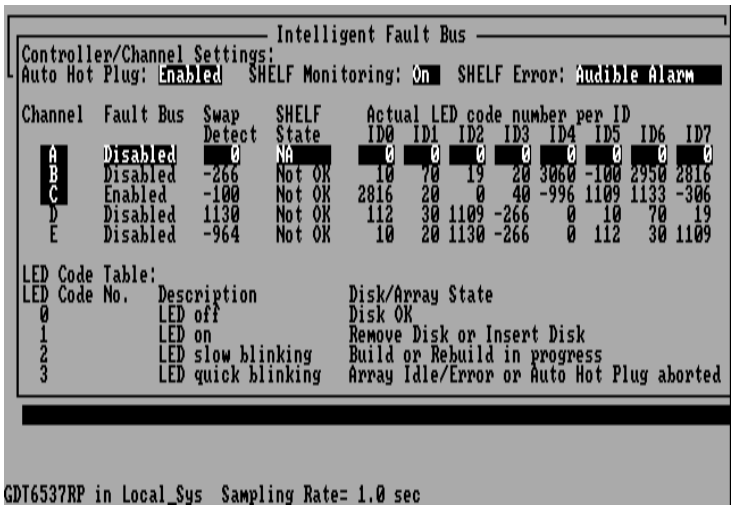

K.3.3 Cache Settings

This submenu displays the current GDT cache settings which can be changed here. The various settings are:

Cache ON	the GDT cache is enabled, that is, all accesses to the Host Drives pass through the GDT cache
Cache OFF	the GDT cache is disabled
Delayed Write ON	Write accesses are delayed, i.e., the write-back cache algorithm is active
Delayed Write OFF	All write accesses are directly transmitted to the Host Drives. If delayed-write is off and the Cache is ON, the GDT cache works exclusively as a read cache.

Note: Best performance is achieved with *Cache ON* and *Delayed Write ON*.

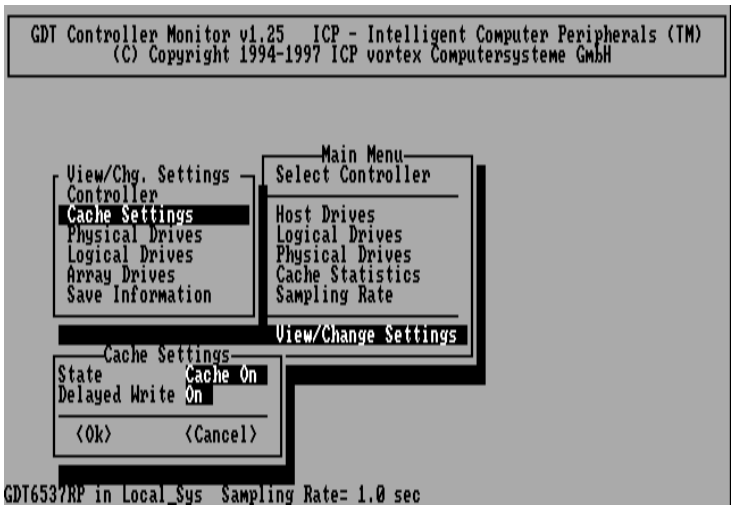

K.3.4 Physical Drives

This option gives a list of all direct access SCSI devices connected to the ICP Controller. Besides information on the GDT SCSI channel, the SCSI-ID, the name/vendor and the gross capacity (1MB = 1024KB), it also shows which SCSI device(s) belong to a given Logical Drive.
When you select a SCSI device in this list and then press <ENTER>, GDTMON gives you further information on the *Drive Settings* of this SCSI device. These settings may be changed.

The Last Status information should always be 0x00000000. After a device failure or other significant events, a different value may be displayed here. Chapter L.7.4 of this manual

gives an excerpt of the SCSI specification with all relevant device Last Status information. This information may be of assistance in further analysis of the problem.

The Grown Defects counter shows the number of media defects which have occurred since the first time the device was operated with a ICP Controller. A specific SCSI hard disk is in a good condition when it has 0 grown defects. When this counter increases, there is definitely something wrong with the device.

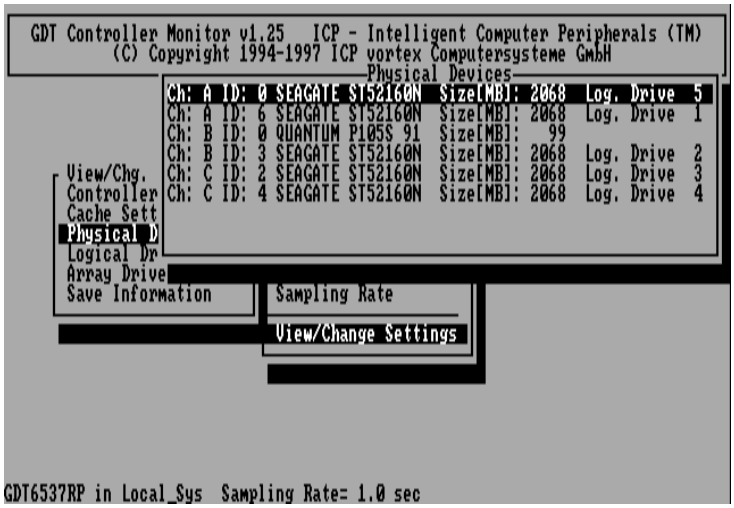

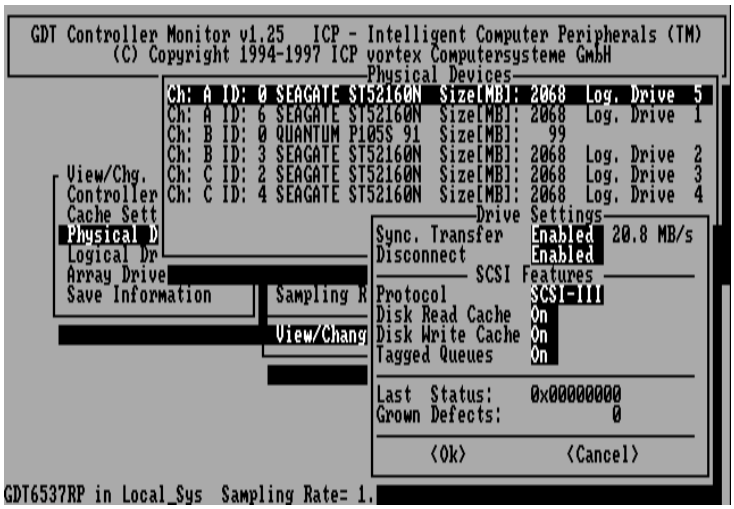

If you select the SEP of a SAF-TE subsystem, GDTMON displays a list of the installed and configured slots in the subsystem

Press <F2> to get the status of the SAF-TE enclosure.

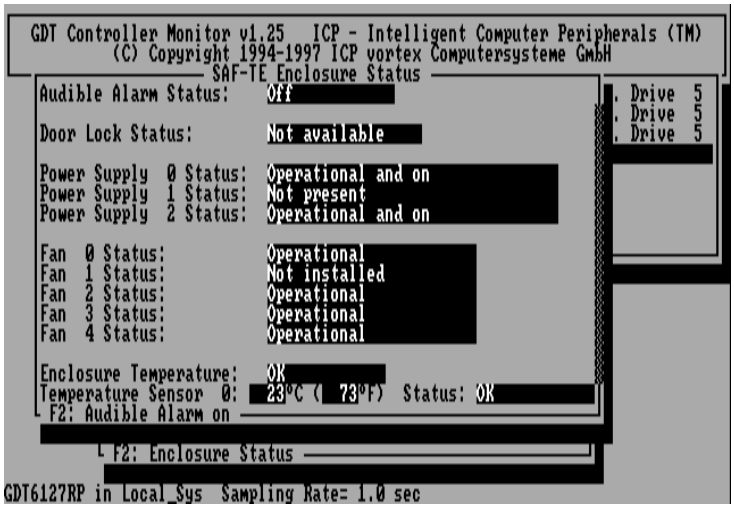

K.3.5 Logical Drives

This command yields a list of the existing Logical Drives. In addition to the Logical Drive numbers, information on the drives' type, state, net capacity and belonging to a given Array Drive / Host Drive is displayed. Press <F2> to obtain further information on a selected Logical Drive.

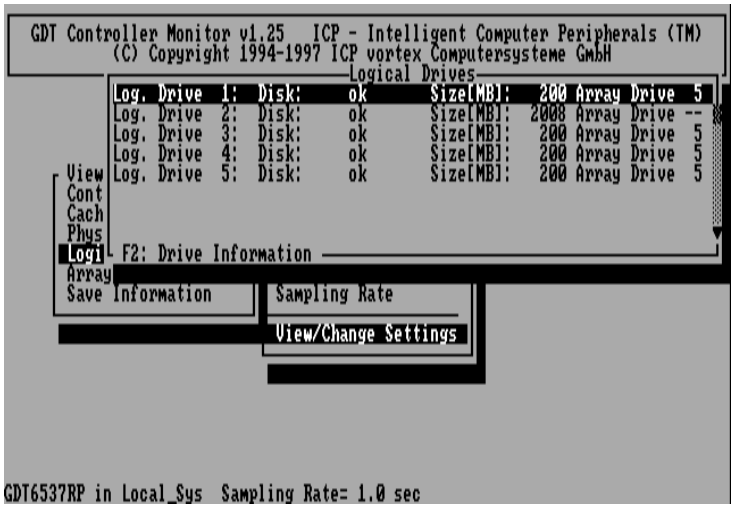

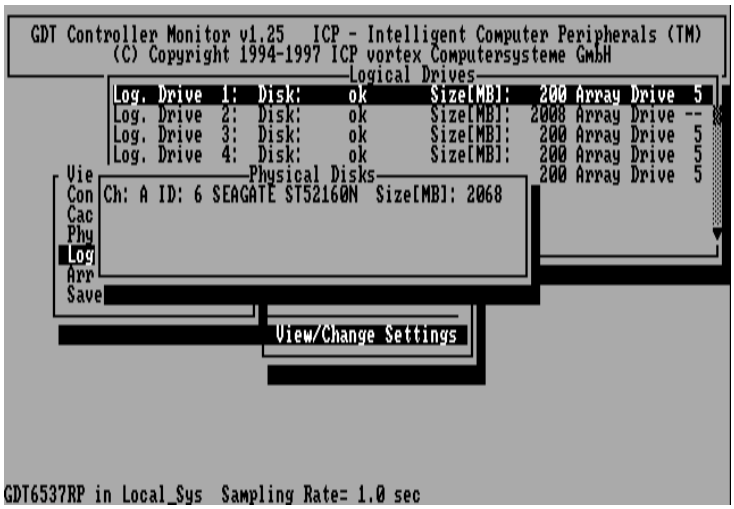

Press <ENTER> to select a Logical Drive. The following options become available:

K.3.5.1 Set Logical Drive Name

Change the name of the Logical Drive. This name was defined within GDTSETUP, either automatically, or manually.

K.3.5.2 Hot Plug: Replace Mirror Drive

If a RAID 1 or RAID 10 Host Drive has already been set up, a defective drive can be replaced (Hot Plug) while the system continues to be fully operational.

There are typically two different applications, where a Hot Plug becomes necessary.

Application 1. The RAID 1/10 Array Drive is in the fault tolerant (both drives are valid: *vv*). It is likely that a drive will fail soon (for example when there is a loud operating noise). As a preventative measure, this drive ought to be exchanged now, that is, in a moment when the Array Drive is still in a ready state and still has redundancy.

Application 2. The Array Drive is no longer fault tolerant, because a drive of the Array Drive has actually failed. The Array Drive is still fully operational, but it does not have redundancy any more. If another drive should fail, too, the Array Drive's functionality is impaired.

Example Session for Application 1

We assume that there is a RAID 1 Array Drive which is fault tolerant. Its state is *vv*, both Logical Drives are valid. After selecting the Array Drive, we choose the Replace Mirror Drive option.

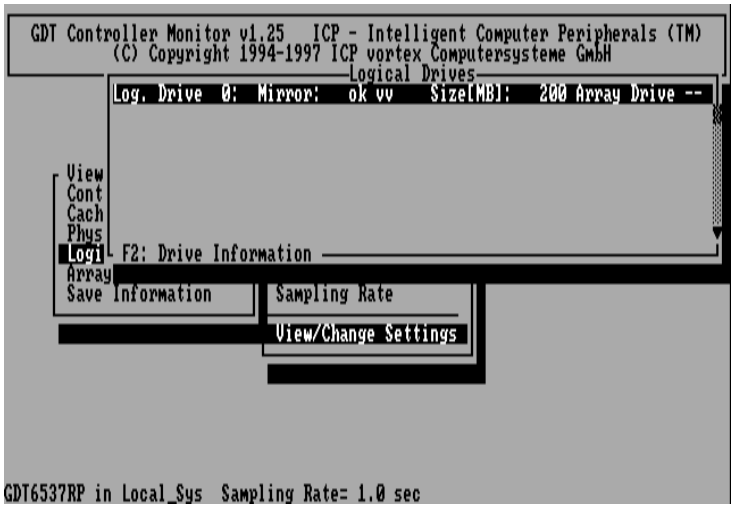

A list is displayed which shows the members of the RAID 1 Array Drive.

The Hot Plug function now displays a list of the positions available for the new drive. Each position is univocally determined by its coordinates (SCSI bus, SCSI ID).Obviously, the new drive can only be assigned to a position which is not occupied by another SCSI device yet, exception made for the position still occupied by the drive to be exchanged.

We choose entry number 1 and press <ENTER>.

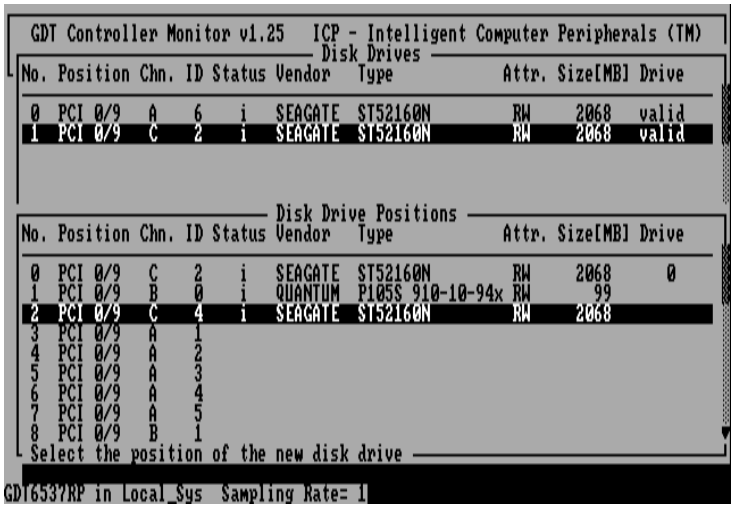

For this example, we select the first position and receive the following message:

The Hot Plug function now informs us that all SCSI devices on the SCSI channel to which the drive to be exchanged is connected, have to be temporarily halted. In addition, it shows which Host Drives are affected by this brief halt. With regard to the new drive, we are given some important information: Required storage capacity, SCSI ID and SCSI bus termination. The <ENTER>-key may only be pressed when the new drive is ready at hand and after having checked its capacity, SCSI ID and SCSI termination (the latter two may have to be

changed). In our example we presume that these preparations have been made, so we press <ENTER> in order to halt the SCSI channel.

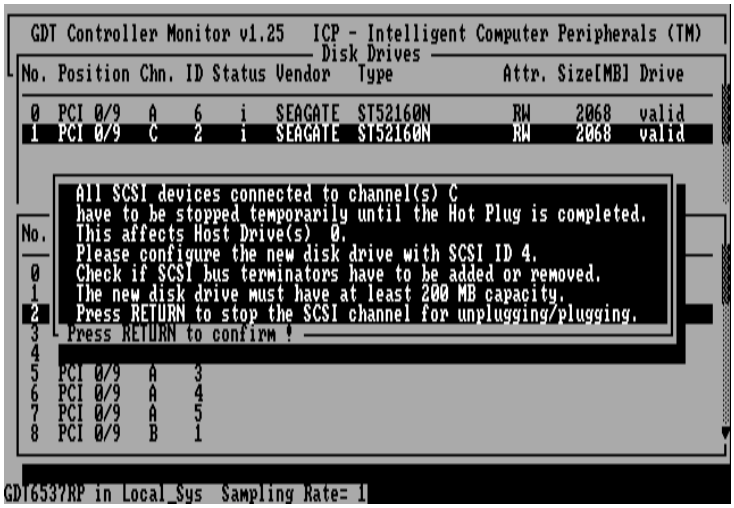

Now we have entered the actual Hot Plug procedure. Disconnect the drive to be exchanged by plugging it off from the SCSI cable first, and then, from the power cable. We immediately connect the new drive to the plugs that are now free, first to the power supply and then to the SCSI bus cable. After having reconnected the new drive properly, press <ENTER>.

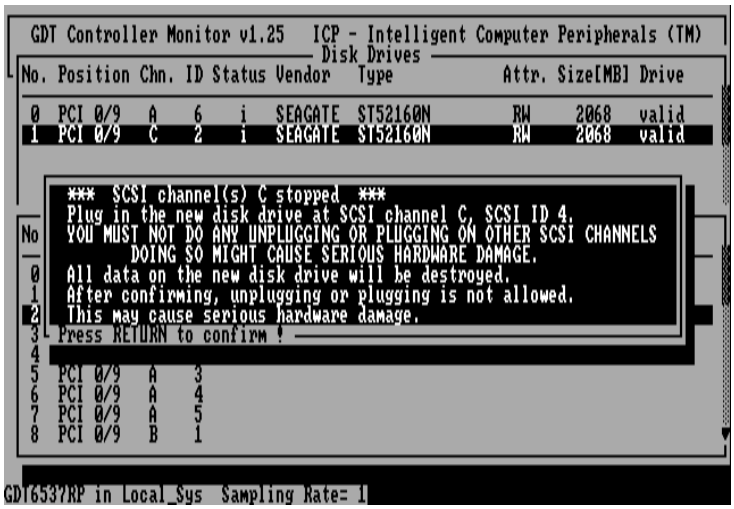

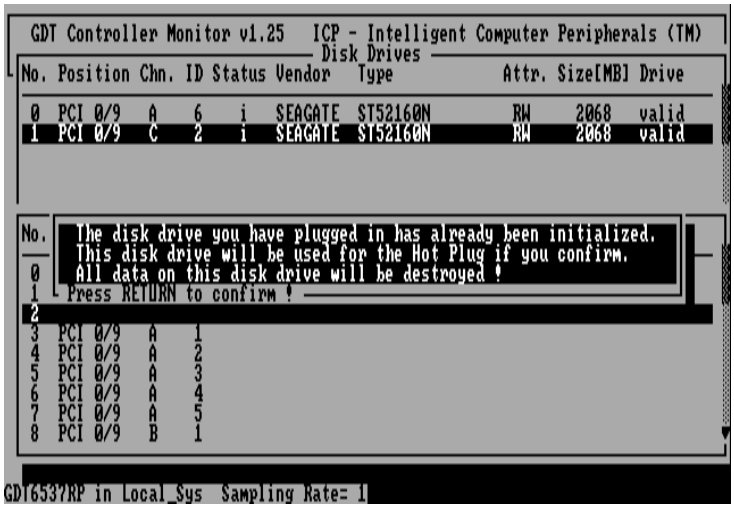

GDTMON recognizes that the new hard disk was already initialized before. Confirmation of this message destroys all data on the selected drive.

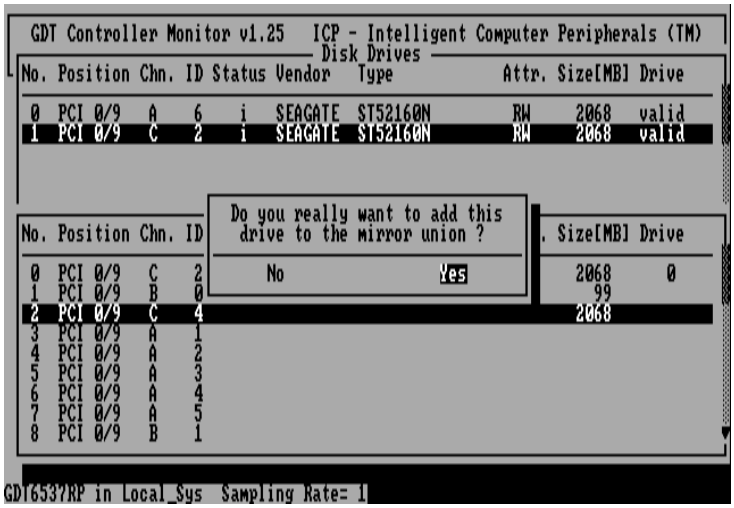

After this confirmation, the Hot Plug is finished successfully. It takes approximately 10 minutes to re-synchronize the data with this sample configuration.

Example Session for Application 2

We assume that there is a RAID 1 Array Drive which is no longer fault tolerant. Its state is *-1/v*, one drive has failed and is therefore no longer accessible on the SCSI bus.
The ICP Controller started beeping. The audible alarm can be disabled within GDTMON by pressing <F2> as soon as the View/Change Settings menu is loaded:

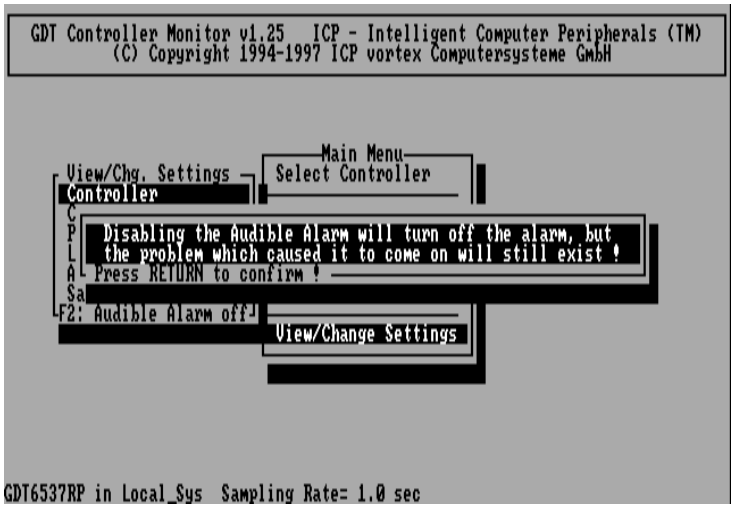

After selecting the Array Drive, we choose the Replace Mirror Drive option.

Before the new drive can be added, the missing drive has to be deleted in the RAID 1 configuration.

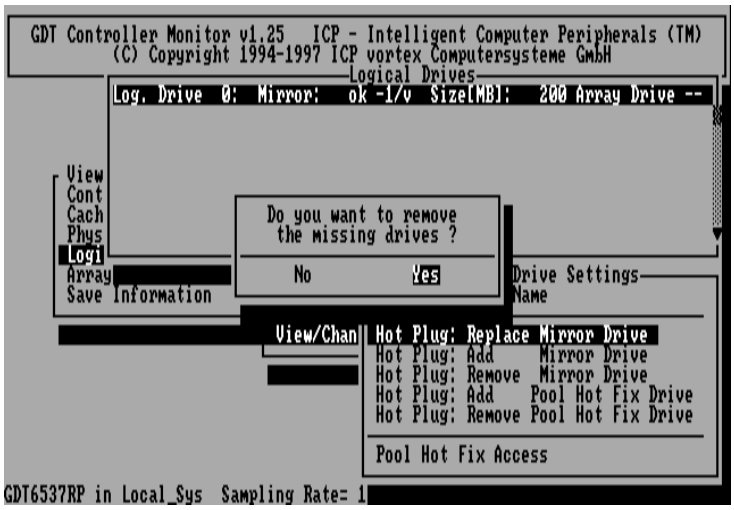

After confirming here with "Yes", you can follow the next paragraph "K.3.5.3 Hot Plug: Add Mirror Drive", to add a new mirror drive to the remaining drive out of the previously failed RAID 1 Array Drive.

K.3.5.3 Hot Plug: Add Mirror Drive

This option allows you to add another Logical Drive as a mirroring drive to another Logical Drive. The new hard disk (configured with a free SCSI-ID and a correct SCSI termination) can be plugged onto the SCSI cable while the system continues to be fully operational. If a spare and suitable hard disk is already connected and not yet assigned to a Logical Drive or Host Drive, it will be displayed in the list of Disk Drive Positions.
There are two cases which make this function very interesting:

1. An existing hard disk should be given 100% redundancy, but there is no time to shut down the system and interrupt the normal operation

2. An existing hard disk has a high probability of failure, because it makes, for example, a strange noise or generates a lot of grown defects, but there is no time to shut down the system and interrupt the normal operation.

Once the new hard disk is added, the data synchronization (mirroring update) is automatically carried out in the background simultaneously with the normal operation.

In this example, there was already a spare hard disk available (on Channel B, SCSI ID 3). This drive must not be a Logical Drive. Otherwise it will not be available for this function.

After selecting the new hard disk, the following message appears:

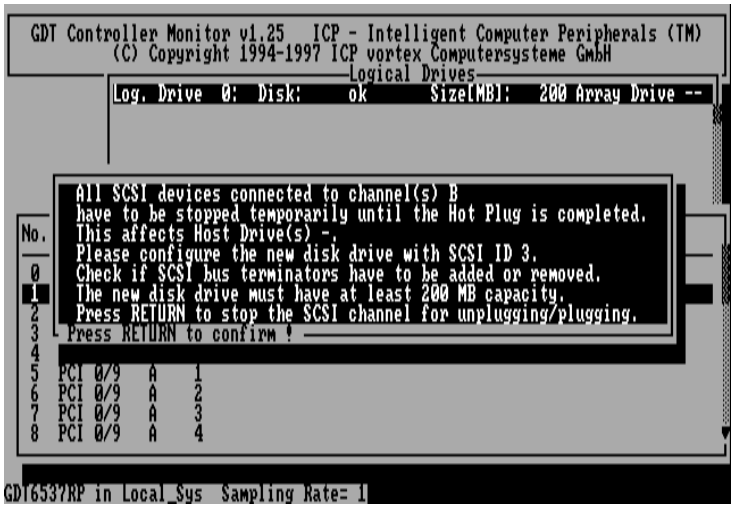

The following message indicates that channel B was stopped for the time of the actual Hot Plug

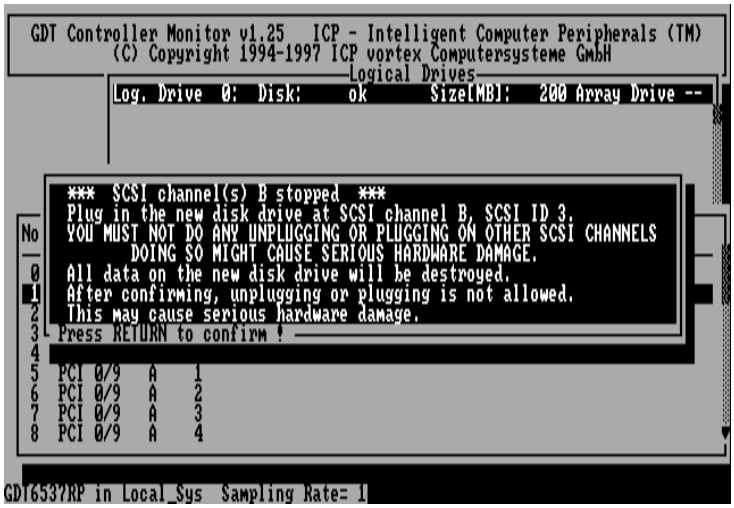

Now, the new hard disk is added as a mirror to the selected Logical Drive. The updated list of available Logical Drives shows the change. The Logical Drive changed its type to *Mirror* and the data on the new hard disk are currently synchronized, indicated through the "*" behind the "v".

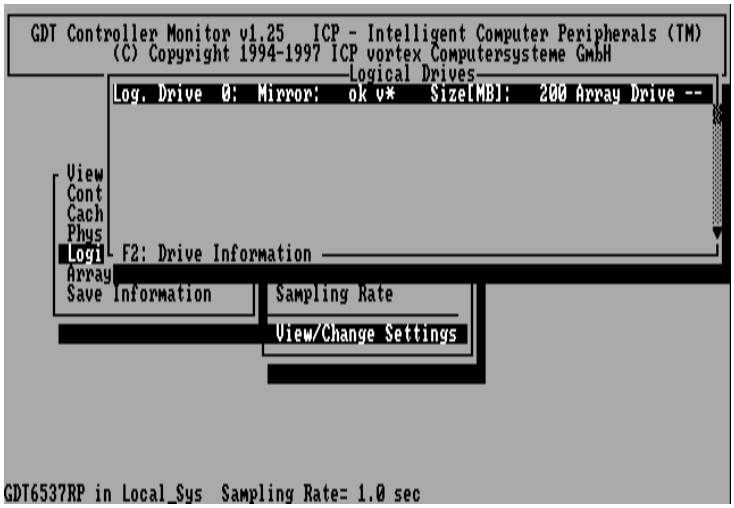

After pressing <F2>, the new structure is displayed.

The entry "invalid" for the second drive means that the data have not yet been (completely) copied from the first drive. After the completion of the synchronization process, this entry changes into "valid".

K.3.5.4 Hot Plug: Remove Mirror Drive

This option allows the removal of a Mirror Drive from a RAID 1 or RAID 10 Array Drive. Once the drive has been removed, the data on the other drive are no longer redundant.

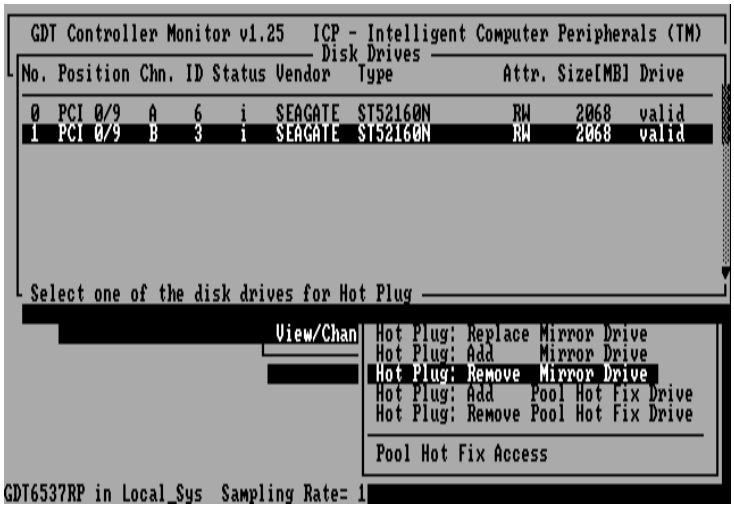

K.3.5.5 Hot Plug: Add Pool Hot Fix Drive

A Pool Hot Fix Drive is a spare drive within the so-called Hot Fix Pool. A drive in a Hot Fix Pool is available for several RAID 1 and RAID 10 Array Drives as a Hot Fix drive. Thus, several Array Drives can share one Hot Fix drive. Of course, once this drive has been used by one of the Array Drives, it is no longer available for the others.

After selecting this option, GDTMON scans the ICP Controller for drives which are suitable

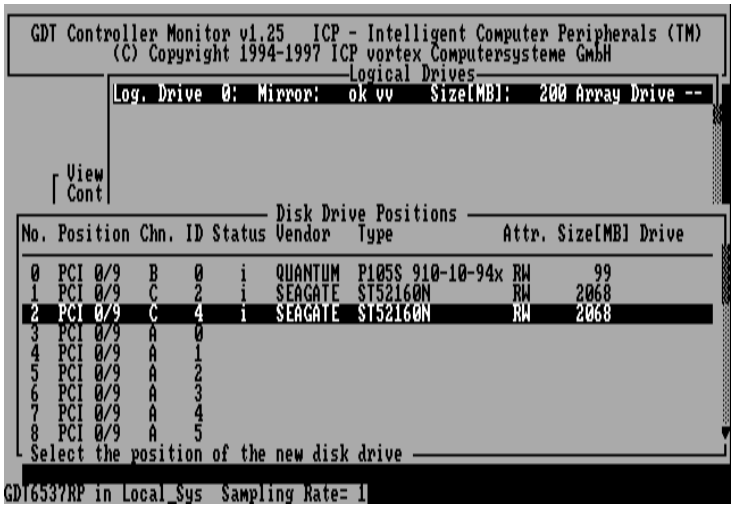

for becoming a Pool Hot Fix drive (i.e. they belong to no Logical Drive), and for free coordinates (SCSI Channel / SCSI ID). We choose Channel C and SCSI ID 4 for the new Hot Fix Drive.

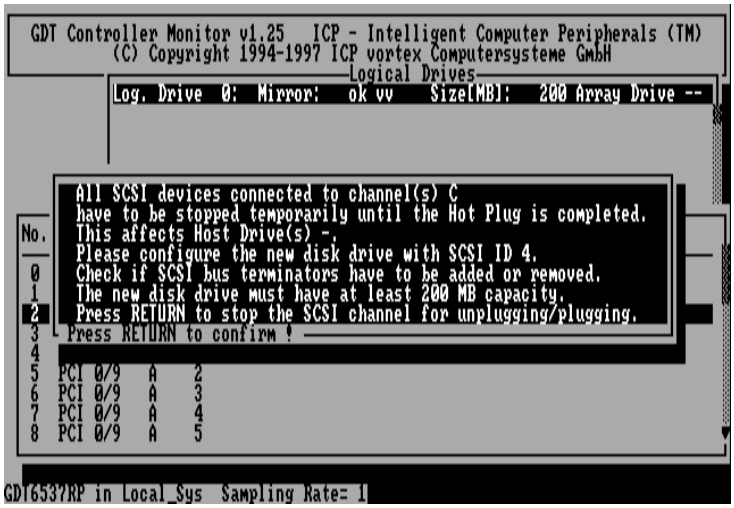

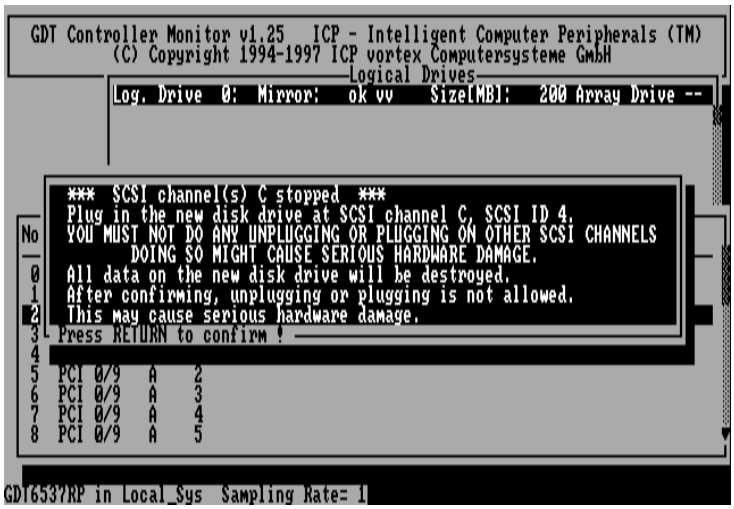

After the completion of this function, the Pool of Hot Fix drives contains a new drive (in our example here, it is the only drive.
To allow a RAID 1 or RAID 10 Array Drive access to the Hot Fix Pool, use the Pool Hot Fix Access menu (K.3.5.7).

K.3.5.6 Hot Plug: Remove Pool Hot Fix Drive

It may become necessary to remove a certain drive from the Hot Fix Pool.

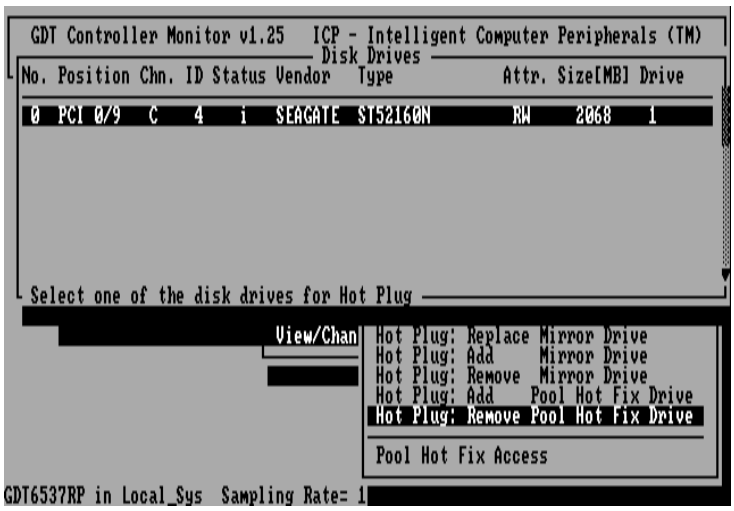

K.3.5.7 Pool Hot Fix Access

This function enables or disables the access of a certain RAID 1 or RAID 10 Array Drive to the Hot Fix Pool.

If the access had been enabled before, you could disable it now.

K.3.6 Array Drives

This command yields a list of the existing RAID 4 and RAID 5 Array Drives. In addition to the Array Drive number, information on the RAID level of the Array Drive: 4 or 5), the state (error, idle, build, ready, fail, expand, rebuild) and the net capacity are displayed. Press <F2> to obtain further information on a selected Array Drive.

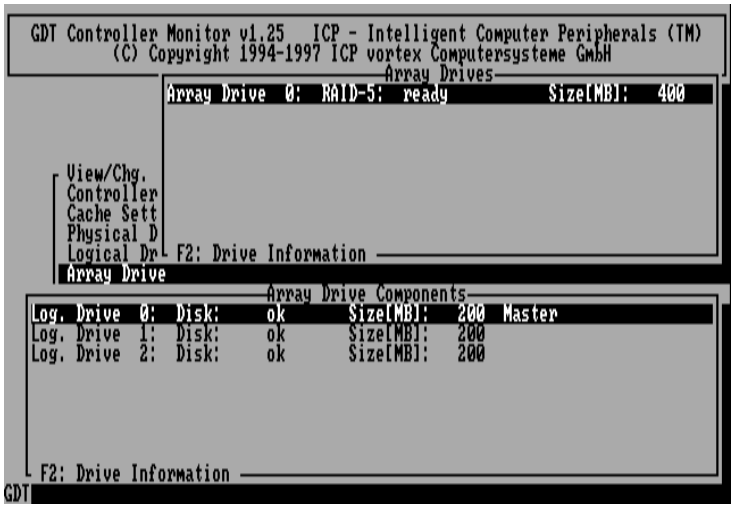

If you press once more <F2>, you get detailed information on the physical hard disk.

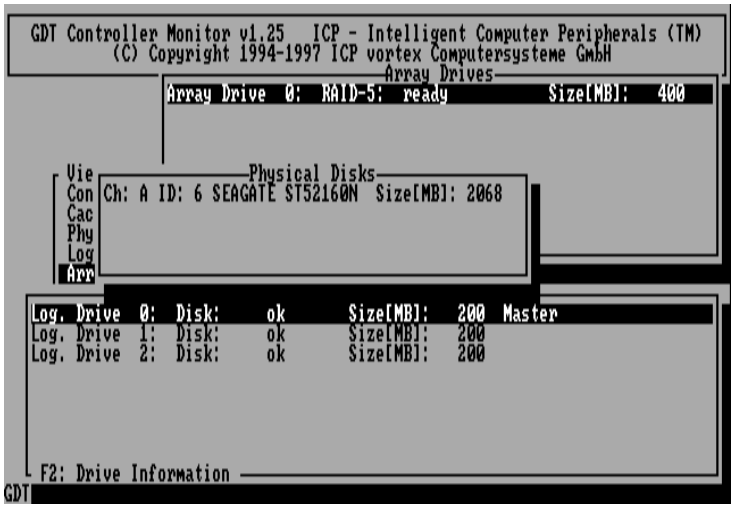

K.3.6.1 Parity Verify

This option verifies online the parity information of the selected RAID 4 or RAID 5 Array Drive. Pressing <ESC> terminates this process. If this option is selected for several Array Drives, the processes are put into a queue and performed one after the other.

K.3.6.2 Parity Recalculate

If the parity verify option reports a parity problem, it is advisable to recalculate the parity of the selected Array Drive anew.

The state of the Array Drive changes into "build/patch", and the build process is started immediately. The word "patch" indicates that the parity of this Array Drive was calculated anew. After this procedure the Array Drive assumes the *ready/patch* sate.

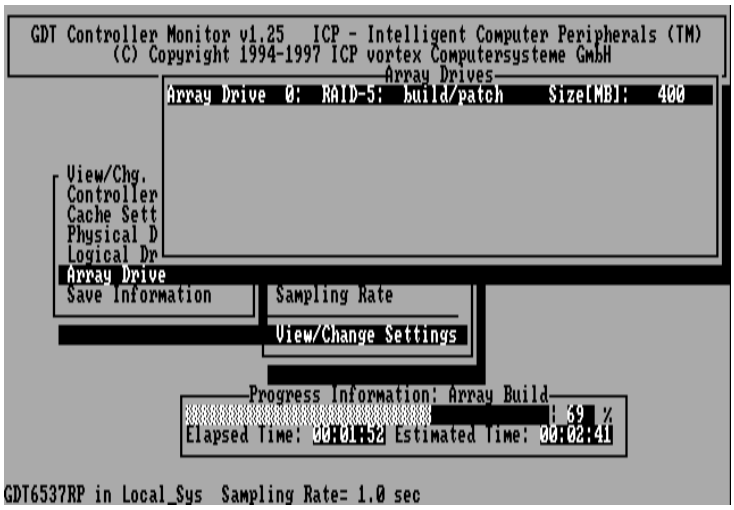

K.3.6.3 Expand Array Drive

There are two fundamental functions which are available within this option:

- Migrate the RAID level of the selected RAID Array Drive
  (RAID 0-> RAID 4 and vice versa, RAID 0 -> RAID 5 and vice versa)
- Expand the capacity of the selected Array Drive by adding one or several new hard disks

Both functions can be selected at the same time. E.g., migrate from RAID 0 to RAID 5 and add a new drive.

To initiate a migration or expansion with a RAID 4/5 Array Drive, the state must be *ready*.
The data on the Array Drive remain intact and are not affected by the expansion.
The additional capacity is introduced as new Host Drive.
If a Logical Drive fails during the expansion, the expansion process continues until the expansion is finished. The Array Drive changes into the *fail* state.

The new capacity is available as a new Host Drive. Windows NT (Tool: Disk Administrator) and Novell NetWare (*Scan for new Devices* and then Tool: Install) allow the online integration of new disk capacity.

Depending on the RAID level the current Array Drive has, selecting a different one here, will cause the Migration of the RAID level of the Array Drive. If you select the same RAID level, the following procedure will expand the capacity of the Array Drive, only.

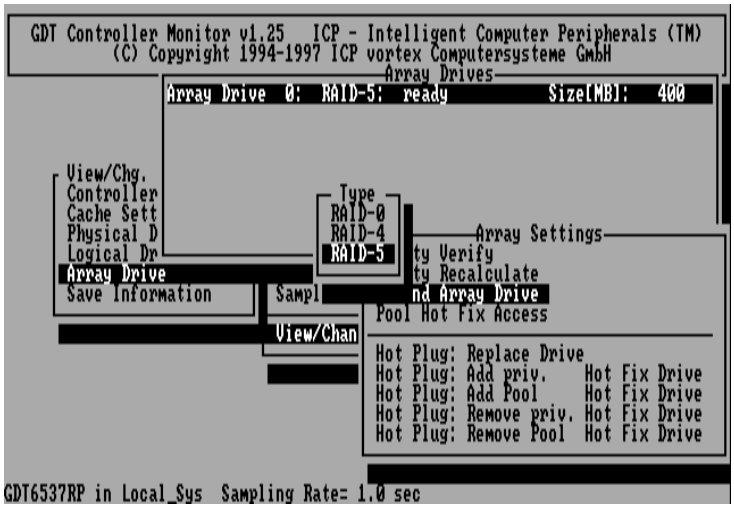

After confirming this request, the SCSI channels are scanned for free positions and already existing available (i.e., not yet assigned to a Logical Drive) hard disks.

For this demo, we select the hard disk on channel A and ID 0.

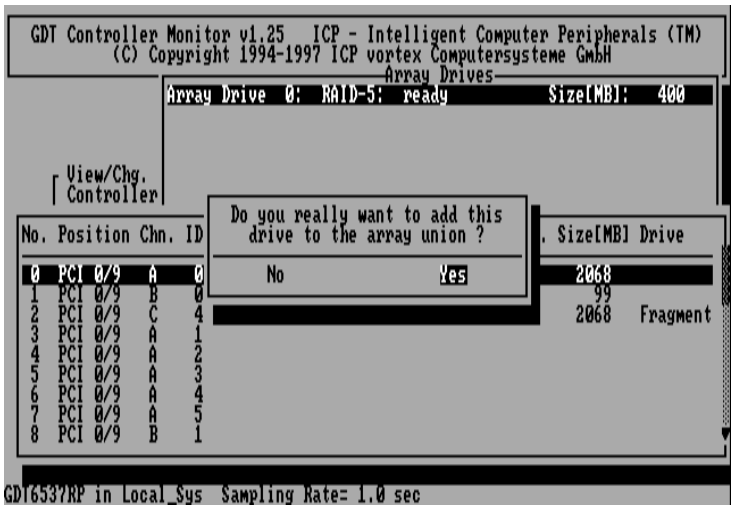

The new drive is built into the Array Drive. According to the Expansion Progress Information this takes approximately 18 minutes. During the expansion the Array Drive's state is *ready/expand*.

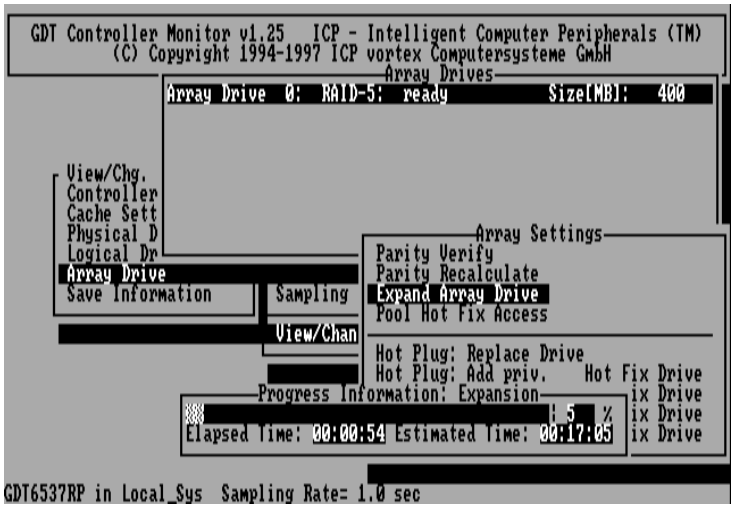

As expected the Array Drive's capacity is now 600MB.

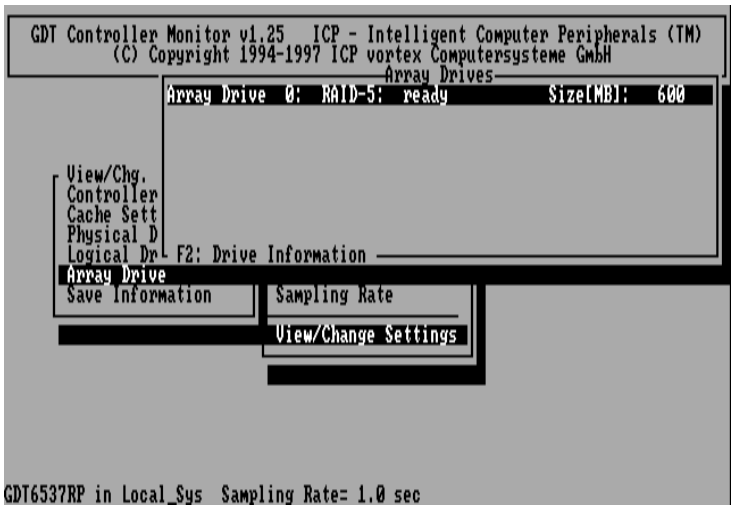

K.3.6.4 Pool Hot Fix Access

This function enables or disables the access of a certain RAID 4 or RAID 5 Array Drive to the Hot Fix Pool.

If the access had been enabled before, you would be able to disable it now.

K.3.6.5 Hot Plug: Replace Drive

In a similar way as was described a few pages before with the RAID 1 Array Drives, this function is designed to replace a defective drive of a RAID 4 or RAID 5 Array Drive, while the system continues to be fully operational.

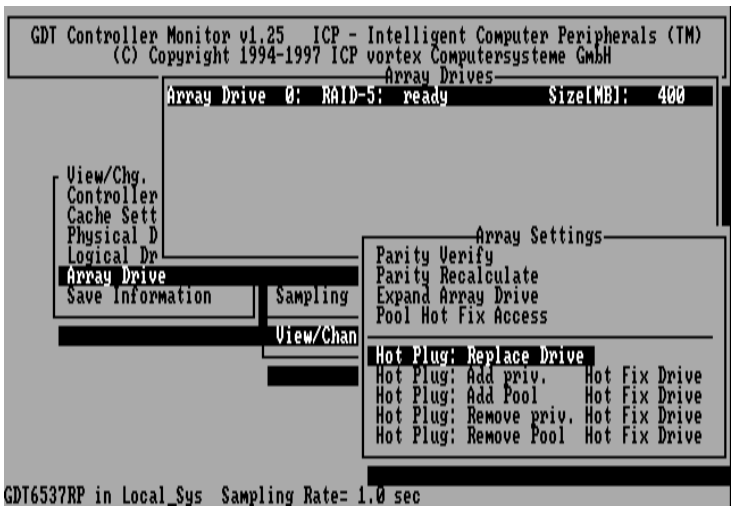

There are typically two different applications where a Hot Plug is necessary.

Application 1. The RAID 4/5 Array Drive is in the *ready* state.
It is likely that a drive will soon fail (for example when there is a loud operating noise). As a preventative measure, this drive ought to be exchanged now, that is, in a moment when the Array Drive is still in a ready state and still has redundancy.

Application 2. The RAID 4/5 Array Drive is no longer fault tolerant (*fail* state), because a drive of the Array Drive has actually failed. The Array Drive is still fully operational, but it does not have redundancy any more. If another drive should also fail, the Array Drive's functionality is impaired.

Example Session for Application 1

We assume that there is a RAID 5 Array Drive which is fault tolerant. Its state is *ready*, all drives are valid. After selecting the Array Drive, we choose the Replace Drive option.

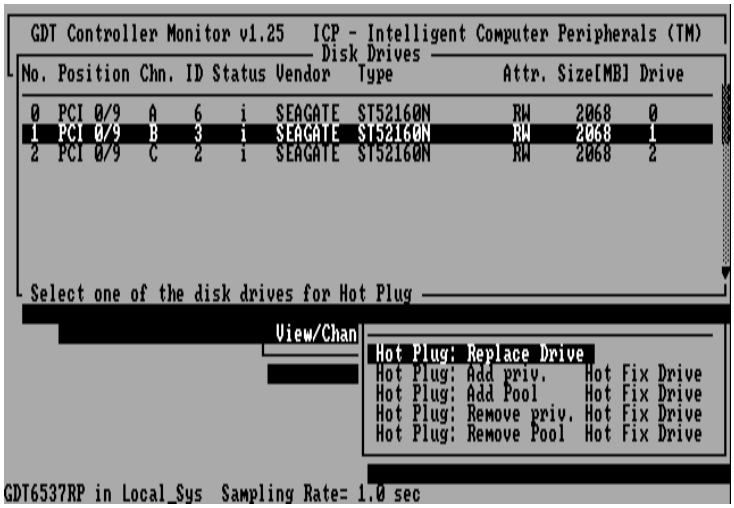

A list of the Array Drive's components is displayed. For our example we choose No.1 for the Hot Plug.
GDTMON scans the ICP Controller's SCSI Channels for drives which are still free (not yet assigned to Logical Drives) and free (i.e., not occupied) SCSI Channels and SCSI IDs.

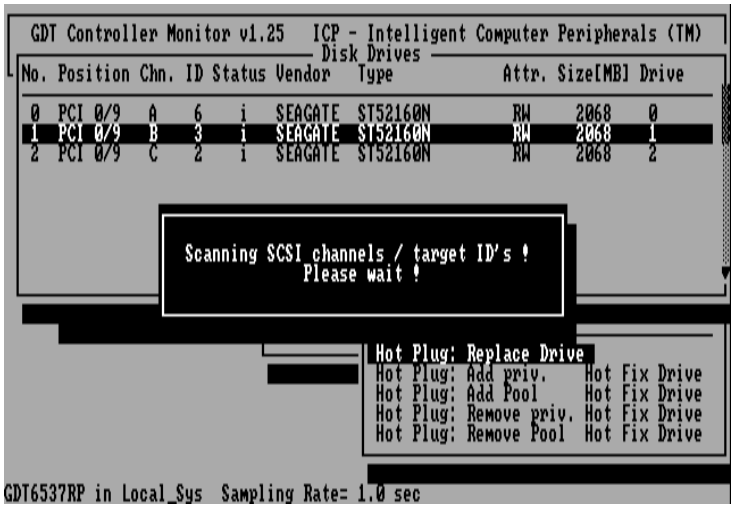

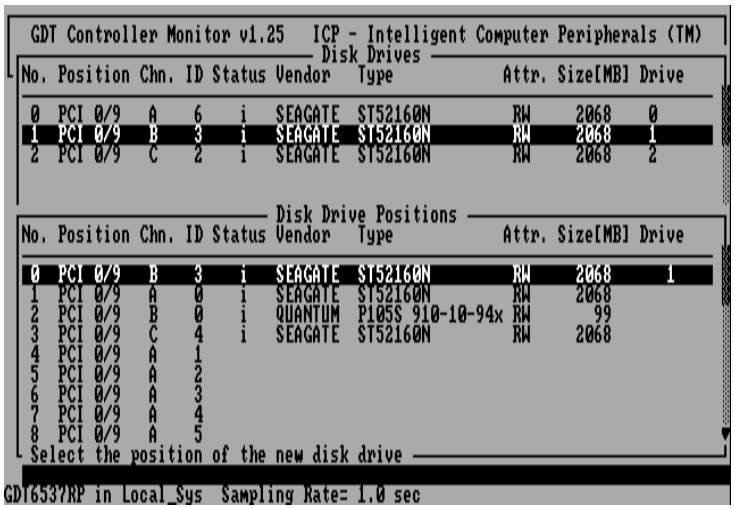

The list of Disk Drive Positions shows us the following:

No.0	This is the position of the drive which should be replaced. Since it is still there, the hard disk's state, vendor, type, attributes, size and Logical Drive number are displayed.
No.1, 2, 3	Free available (i.e., not yet assigned to a Logical Drive) drives. 1 and 3 could also be used for the replacement.
No.4 to No.17	These are free (i.e., not occupied) plugging positions for the new hard disk (On the above screen you can't see positions No.9 to No.17, but you can scroll the list with the Page-Down/Up keys).

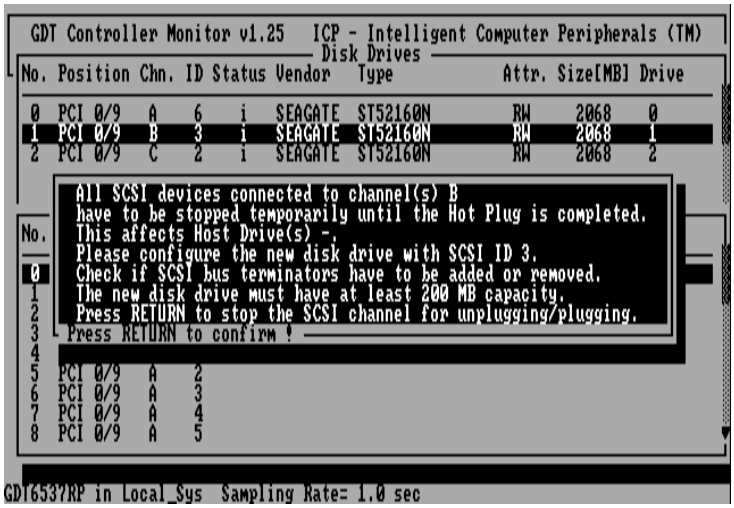

For our example we now take the new hard disk (which must have a capacity equal or larger than 200MB) set it to SCSI ID 3 and observe the SCSI termination.

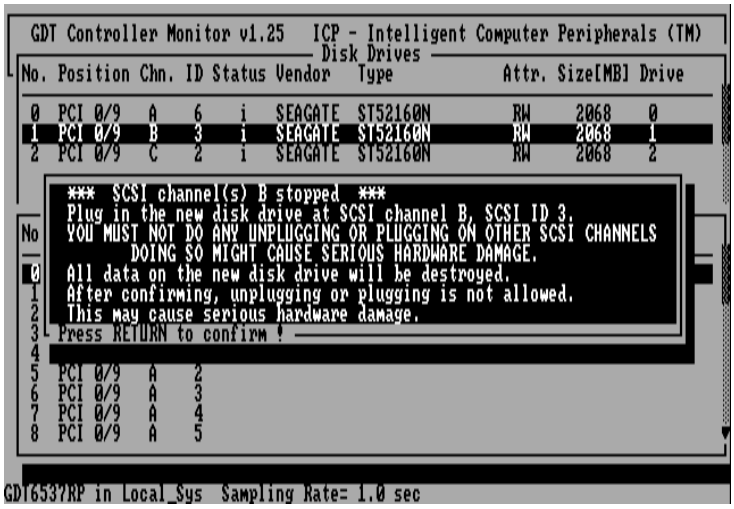

After this message we can unplug the old drive and plug in the new one again and confirm this procedure.

If everything was OK, GDTMON displays the following message:

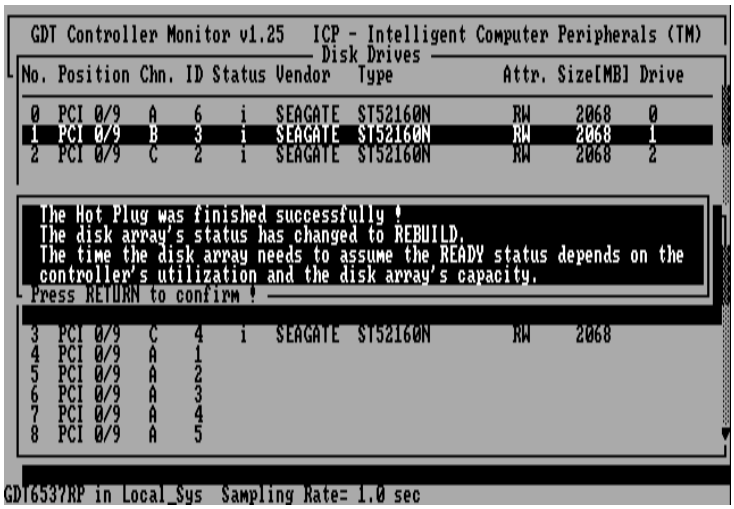

If the new drive, which we have plugged in just before, had contained data from a previous operation with a ICP Controller, GDTMON would have reported this.

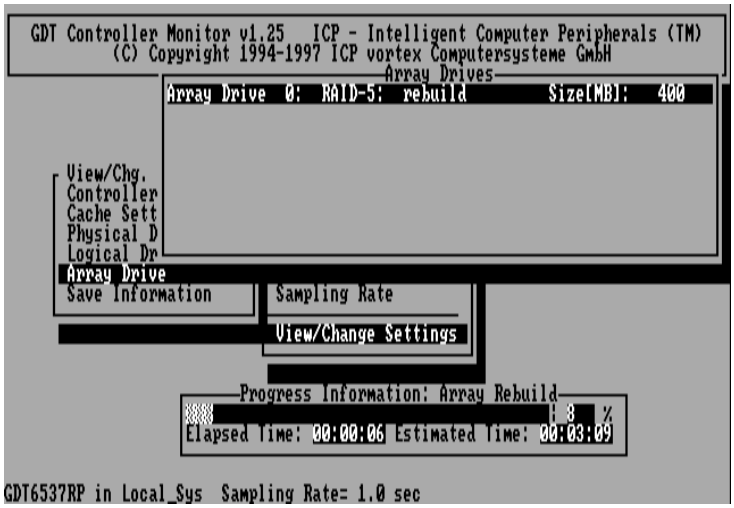

Example Session for Application 2

We assume that there is a RAID 5 Array Drive where one drive has failed. Its state is *fail*. After selecting the Array Drive, we choose the Replace Drive option.

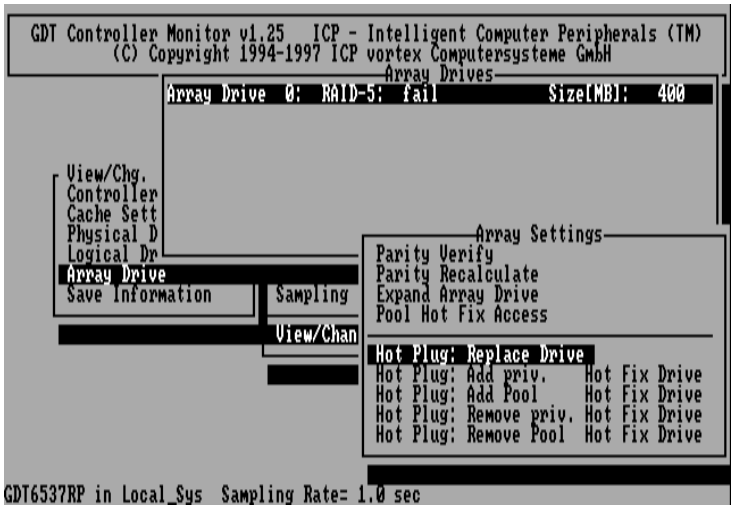

GDTMON shows the failed drive (No.1),

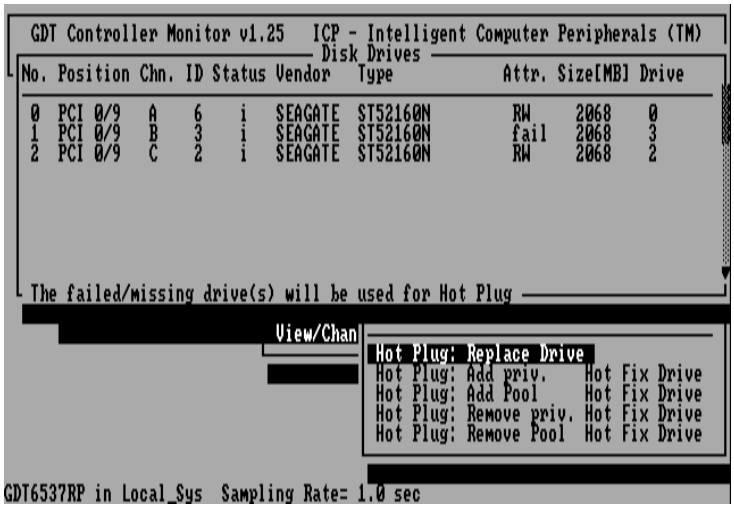

After pressing the <ENTER>-key, GDTMON scans the ICP Controller for free plugging positions:

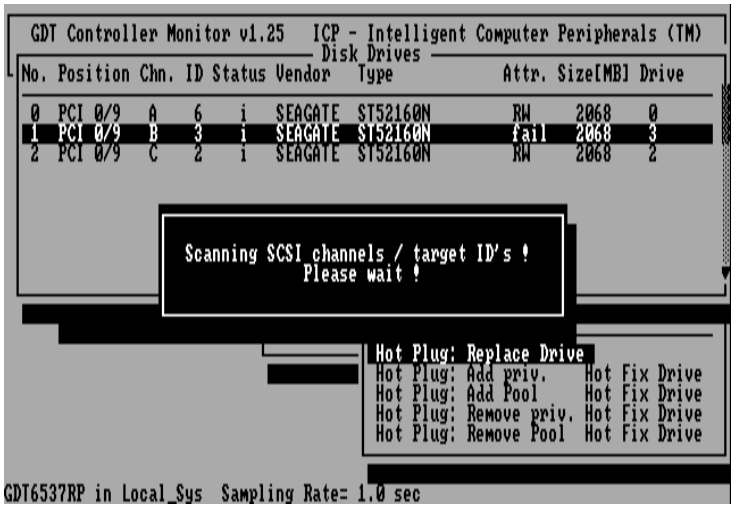

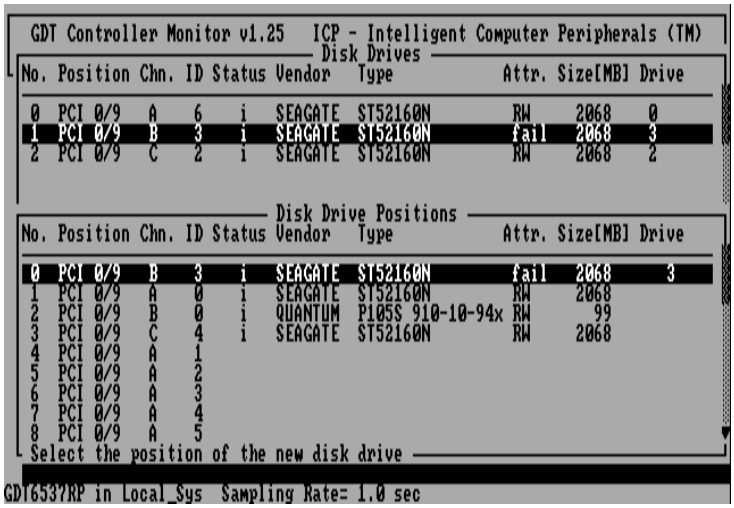

The list of Disk Drive Positions shows us the following:

No.0	This is the position of the drive which should be replaced. Since it is still there, but defective, the hard disk's state, vendor, type, attributes, size and Logical Drive number are displayed.
No.1, 2, 3	Free available (i.e., not yet assigned to a Logical Drive) drives. 1 and 3 could also be used for the replacement.
No.4 to No.17	These are free (i.e., not occupied) plugging positions for the new hard disk (On the above screen you can't see positions No.9 to No.17, but you can scroll the list with the Page-Down/Up keys).

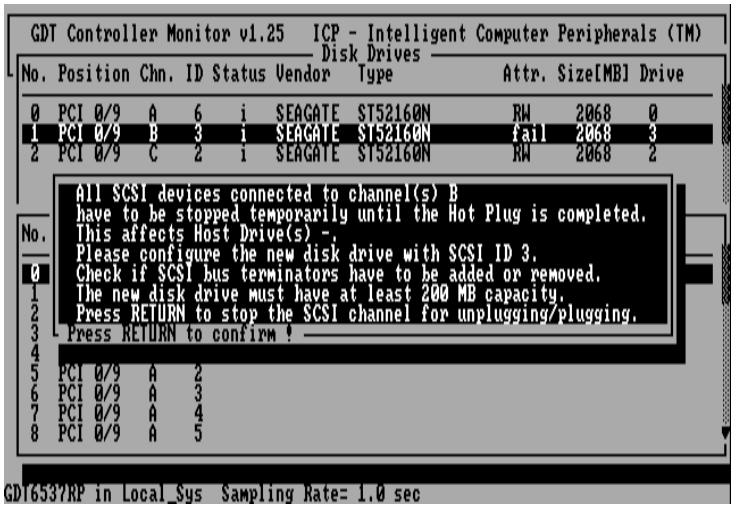

For our example we now take the new hard disk (which must have a capacity equal to or larger than 200MB) set it to SCSI ID 3 and observe the SCSI termination.

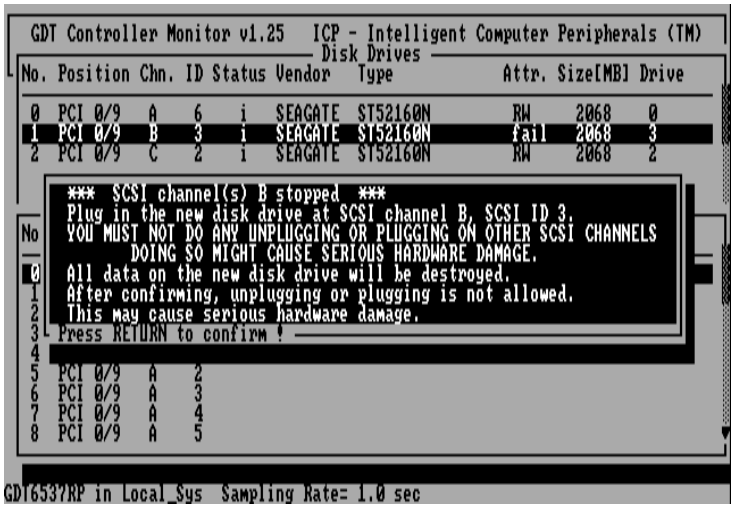

After this message we can unplug the old drive and plug the new one in again and confirm this procedure.

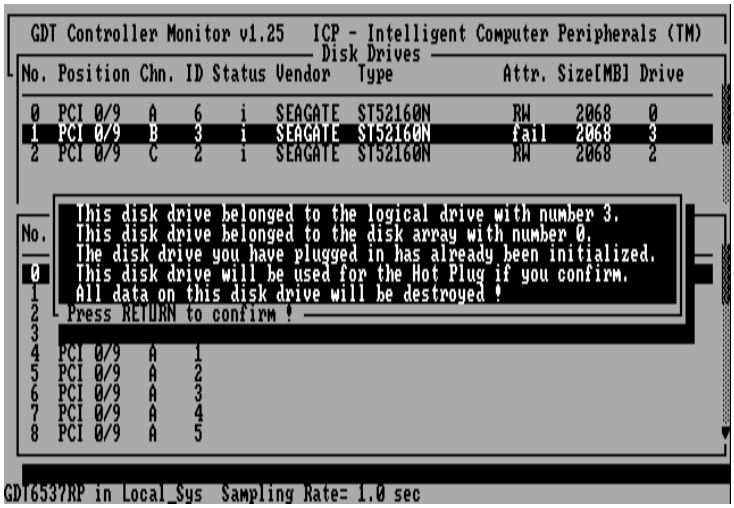

GDTMON has detected data on the new drive (i.e., it was already used as a Logical Drive with a ICP Controller).

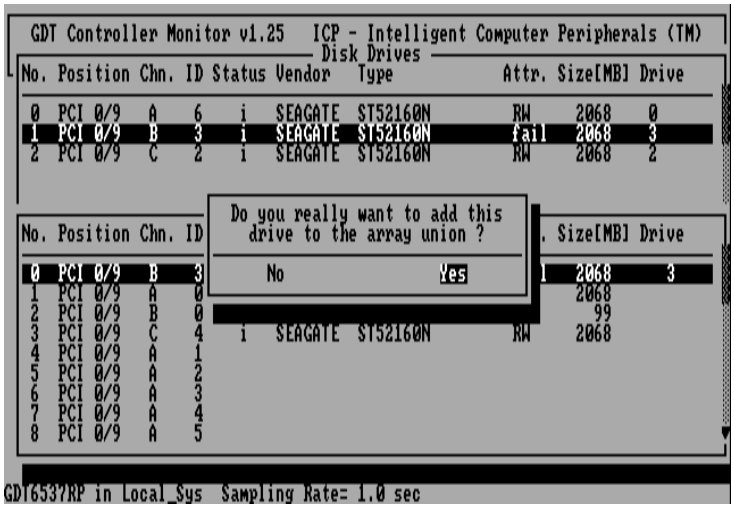

This confirmation deletes all data on the new drive and prepares it for the Array Drive.

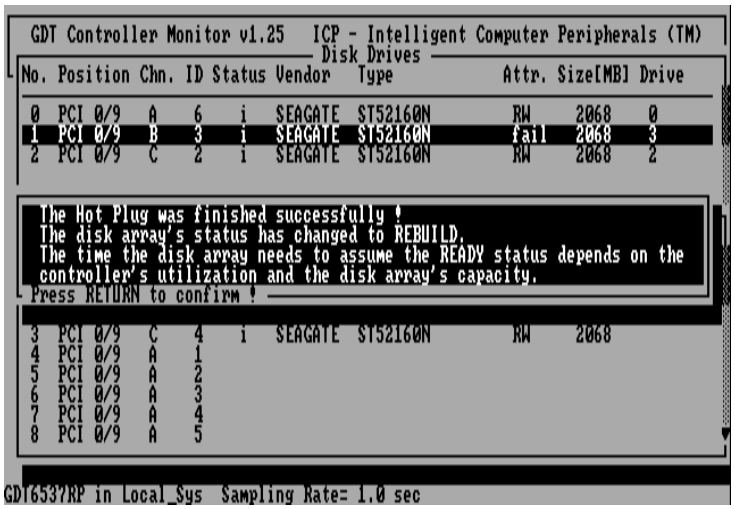

The Array Drive changes its state to *rebuild*. This means that the ICP Controller rebuilds the original data on the new drive.

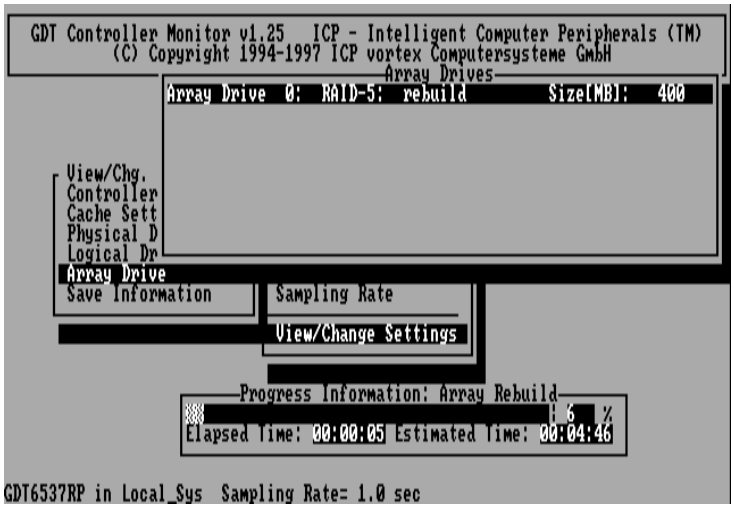

K.3.6.6 Hot Plug: Add private Hot Fix Drive

This function allows you to add a Hot Fix Drive to an existing RAID 4 / RAID 5 Array Drive. "Private" means that this Hot Fix Drive is only available for the selected Array Drive and cannot be accessed from other Array Drives.

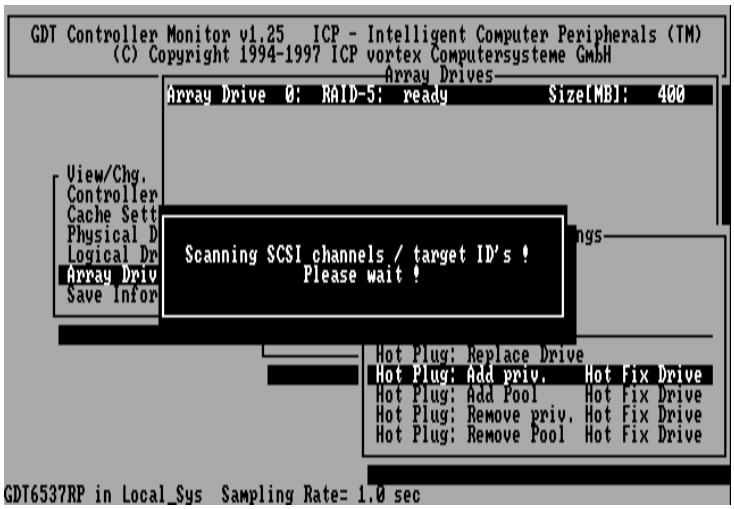

After selecting this option GDTMON scans the ICP Controller for free positions where the new Hot Fix Drive can be plugged in.

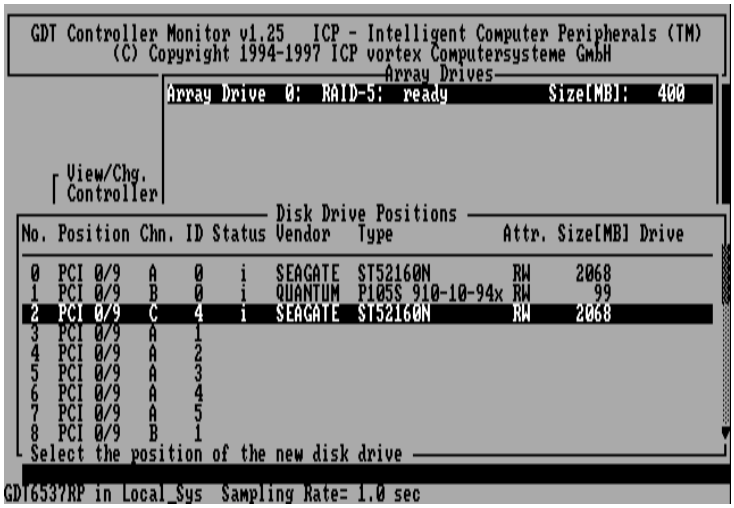

For our example, we choose the drive on Channel C and SCSI ID 4. (If we would plug in now a new drive we would have to set the SCSI ID to 4 and make sure that the SCSI termination is set properly).

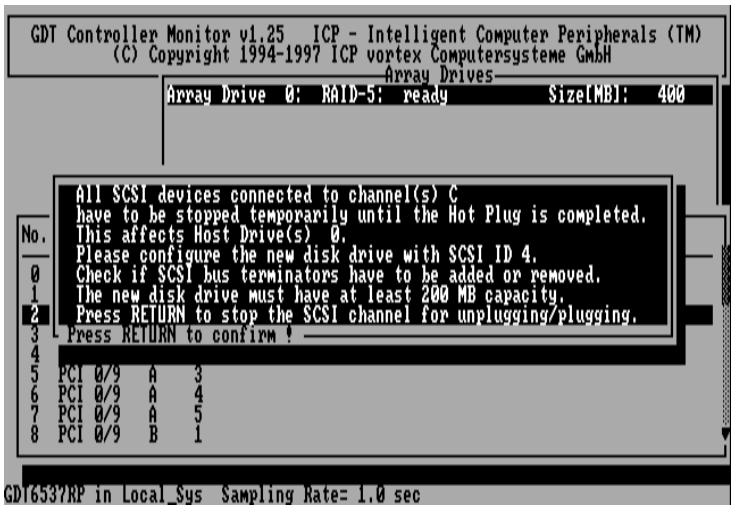

GDTMON adds the new drive to the selected Array Drive.

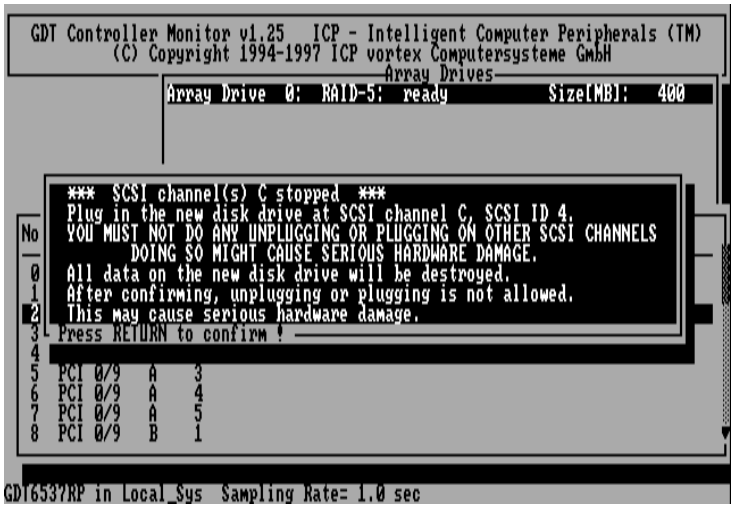

If we now look at the Array Drive's structure (press <F2>), we can see the new drive added as a Hot Fix Drive to the Array Drive.

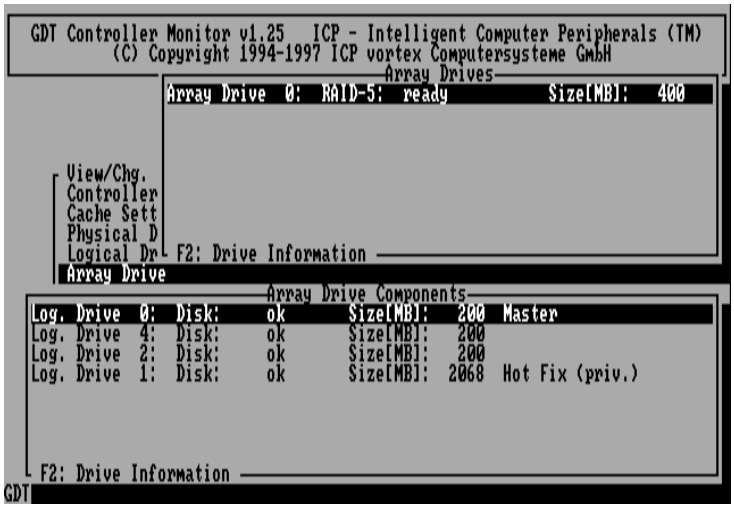

K.3.6.7 Hot Plug: Add Pool Hot Fix Drive

A Pool Hot Fix Drive is a spare drive within the so-called Hot Fix Pool. A drive in a Hot Fix Pool is available for several RAID 4/5 Array Drives as a Hot Fix drive. Thus, several Array

Drives can share one Hot Fix drive. Of course, once this drive has been used by one of the Array Drives, it is no longer available for the others.

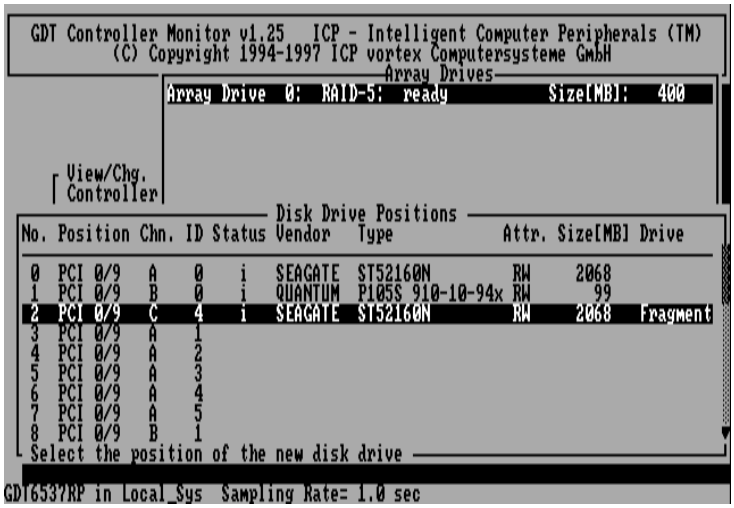

We select SCSI Channel C and ID4 for the new Pool Hot Fix Drive.

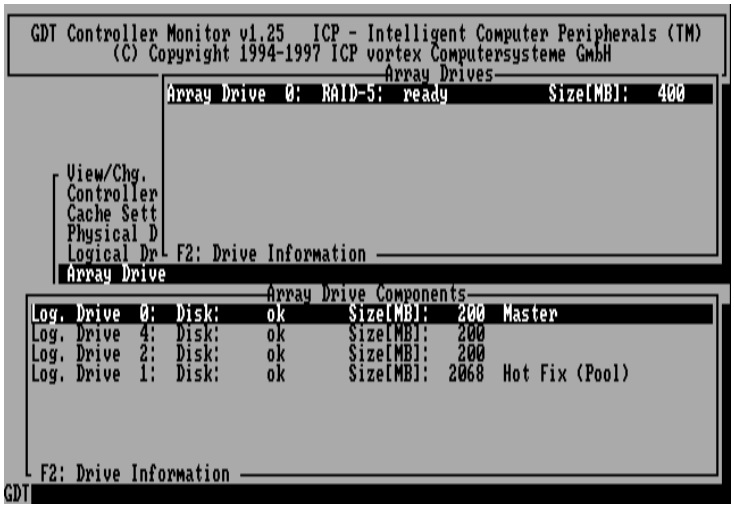

You may disable or enable the access of a certain RAID 4/5 Array Drive to the pool of Hot Fix Drives with the option "Pool Hot Fix Access" (see K.3.6.4).

K.3.6.8 Hot Plug: Remove Private Hot Fix Drive

This function is used, if you want to remove a private Hot Fix Drive from an Array Drive.

K.3.6.9 Hot Plug: Remove Pool Hot Fix Drive

In a similar way as with the "Remove private Hot Fix Drive" function, here you can remove a Hot Fix Drive from the Hot Fix Pool. A possible reason for this could be that you want to add it as a private Hot Fix Drive to an Array Drive.

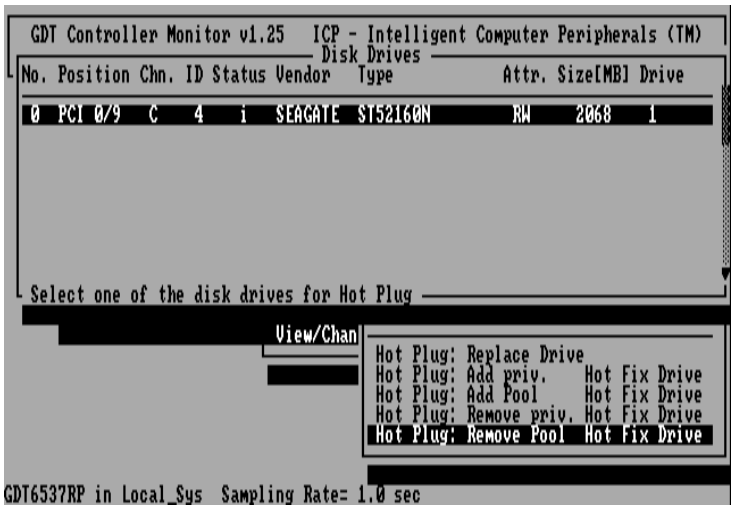

K.3.7 Save Information

The *Save Information* option gives you the possibility to save the configuration information regarding the selected ICP Controller and its SCSI devices in an ASCII-file. This may help if you require support and is also good for your system documentation.

This page is intentionally left blank.