

Addendum

ICP Controllers In Cluster Systems

4th Edition

© Copyright 1992-2001

ICP vortex Computersysteme GmbH
Konrad-Zuse-Str. 9
74172 Neckarsulm - Germany

ICP vortex Corporation
4001 E. Broadway Rd., B-20
Phoenix, AZ 85040, USA

All Rights and Changes reserved.





Contents

Chapter A	General Information
Chapter B	Installation Notes
Chapter C	ICP RAID Console
Chapter D	ICP RAID Navigator
Chapter E	Appendix

For information on the installation of the ICP Disk Array Controller and the configuration of Disk Arrays, please refer to the "User's Manual" or "Quick Installation Guide" for your ICP Controller. These documents may also be found on the ICP System CDROM.

For a successful installation of Microsoft® Cluster Server, Novell Orion Cluster or Watchdog Cluster of Apptime Inc., it is very important that the person who carries out the installation has detailed knowledge and experience in the areas Operating System Installation and Maintenance, Clustering and ICP Disk Array Controllers.

Limited Warranty

ICP vortex Corporation ("ICP vortex") guarantees that this product is free from defects in material and workmanship. Subject to the conditions and limitations set forth below, ICP vortex will, at its own

option, either repair or replace any part of this product which proves to be defective by reasons of improper workmanship or materials. Parts used to repair products or replacement products will be provided by ICP vortex on an exchange basis, and will be either new or refurbished to be functionally equivalent to new.

This warranty does not cover any damage to this product, which results from accident, abuse, misuse, natural or personal disaster, Acts of God, or any unauthorized disassembly, repair or modification. The duration of this warranty is two years from the date of original retail purchase.

Warranty Claim Requirements

To obtain warranty service, return the defective product, freight prepaid and insured, to your local authorized ICP vortex dealer or distributor, or to ICP vortex Corporation, 4001 E. Broadway, B-20, Phoenix, AZ 85040. Please note the following: You must include the product serial number, and a detailed description of the problem you are experiencing. You must also include proof of the date of original retail purchase as evidence that the product is within the warranty period.

If you need to return the product to ICP vortex, you must first obtain a Return Material Authorization (RMA) number by calling ICP vortex Corporation at 602-414-0414. This RMA number must be displayed on the outside of your package. Products must be properly packaged to prevent damage in transit. ICP vortex accepts no responsibility for products which are damaged on arrival due to poor freight service.

Disclaimers

The foregoing is the complete warranty for ICP vortex products and supersedes all other warranties and representations, whether written or oral. Except as expressly set forth above, no other warranties are made with respect to ICP vortex products. ICP vortex expressly disclaims all warranties not stated herein, including, to the extent permitted by applicable law, any implied warranty of merchantability or fitness for a particular purpose. In no event will ICP vortex be liable to the purchaser, or to any user of the ICP vortex product, for any data loss, data corruption, damages, expenses, lost revenues, lost savings, lost profits, or any other incidental or consequential damages arising from the purchase, use or inability to use the ICP vortex product, even if ICP vortex has been advised of the possibility of such damages.

ICP vortex is not liable for, and does not cover under warranty, any costs associated with servicing and/or installation of ICP vortex products.

This manual has been validated and reviewed for accuracy. The sets of instructions and descriptions were accurate for ICP Disk Array Controllers at the time of this manual's production. However,

succeeding Controllers, software and manuals are subject to change without notification. Therefore, ICP vortex assumes no liability for damages incurred directly or indirectly from errors, omissions or

discrepancies between the Controller, software and the manual.



Pick up the phone



**if you need technical support
and dial the numbers:**

For Europe: +49-(0)7132-9620-900

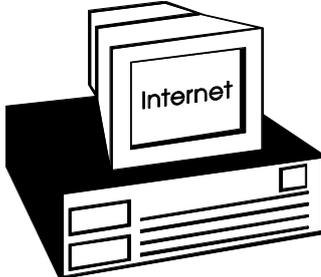
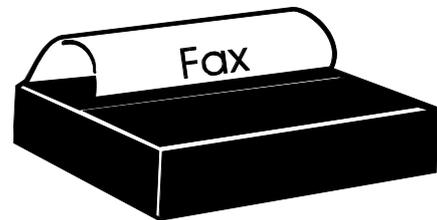
For the USA: 602-414-0414

or send us a FAX:

For Europe:

+49-(0)7132-9620-400

For the USA: 602-414-0444



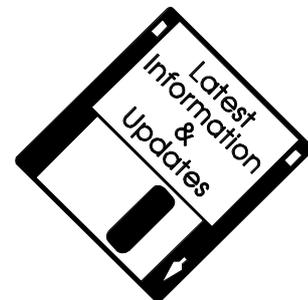
or send us an E-Mail:

For Europe: support@vortex.de

For the USA: support@icp-vortex.com

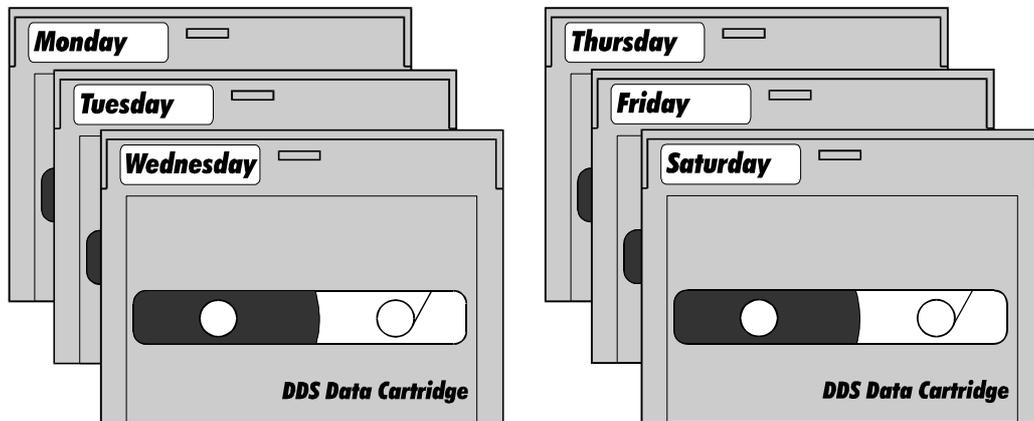
or check our Website:

<http://www.icp-vortex.com>



Important Note

Using modern RAID Systems significantly increases data security and availability. *Under no circumstances* does it relieve you from a careful and daily backup on tape or a similar backup media. This is the only method to protect your valuable data against total loss (e.g., through fire or theft), accidental deletion, or any other destroying impacts.



Many Thanks to all my Friends

Monika & Wolfgang (the grandmasters)
AnnDee, Lois, Valerie, Frank, Ken and Andreas (the Phoenix Crew)
Achim, Dieter, Günter, Hooshiar, Jörg, Norbert, Otto, Ralph, Sam, Steffen, Winfried
Alfred (AB, "We need clustering. I say we have it")
Andreas (AK, or "Kopf nur mit ö")
Michael (Mipf, "where is my CPU ?")
Jürgen (Jogo, "Hi, is Jurgen there ?")
Jürgen (JB, "diesbezüglich & hinsichtlich or probably")
Markus (Malu, "Luuuuu..."), Uwe

All the fantastic "rest" of this incredible company.

It is not only a pleasure to work here, it is a passion.

FCC Compliance Statement

Information for the User

NOTE: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in residential installations. This equipment generates, uses, and can radiate radio frequency energy, and if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorientate or relocate the receiving antenna.
- Increase the separation between the equipment and the receiver.
- Plug the equipment into an outlet on a circuit different from that to which the receiver is powered.
- If necessary, consult the dealer or an experienced radio/T.V. technician for additional suggestions.

The use of a non-shielded interface cable with the referenced device is prohibited. Changes or modifications not expressly approved by ICP vortex Computersysteme GmbH could void the authority to operate the equipment.

Table of Contents

A. INTRODUCTION	12
A.1 ICP Controllers in a Cluster Configuration	13
A.2 Copyrights, Patents	14
A.3 Software License Agreement	15
A.4 ICP vortex Accessories for SCSI Controllers.....	16
A.4.1 SCSI Wonder.....	16
A.4.2 SCSI Wonder Overview	18
A.4.3 Cables, Terminators, Brackets	19
A.4.4 SCSI Wonder to set up Cluster-Systems.....	20
A.5 ICP vortex Accessories for Fibre Channel Controllers.....	23
A.5.1 ICP Fibre Hub 1016.....	23
A.5.3 Sample Configurations	27
B. INSTALLATION NOTES	30
B.1 Hardware Installation	30
B.2 Using Wide/Ultra2 and Ultra 160 SCSI ICP Controllers.....	30
B.4 Releasing the Cluster Channels.....	31
B.5 The Microsoft® Cluster Server Concept.....	32
B.6 Notes on Novell Netware Orion Clustering	34
B.6.1 Hardware Requirements and Settings	34
B.6.2 Install Cluster Array Drives.....	34
B.6.3 Installation of Netware OS and Cluster Services	35
B.6.4 The Netware Client	35
B.6.5 Information on the Active-Active Failover Scenario	36
B.6.6 The ICP Netware Driver and the Tool ICPCON.NLM.....	36
B.6.7 Restrictions.....	36
B.7 Notes on the ICP RAID Navigator (ICPRNAV).....	37
B.8 Notes on ICP RAID Console (ICPCON)	40
C. THE PROGRAM ICPCON	42
C.1 Loading ICPCON.....	42
C.1.1 Loading the ICPCON Program Under NetWare	42
C.1.2 Loading the ICPCON Program Under Solaris 7	43
C.1.3 Loading the ICPCON Program Under Windows NT / 2000	43
C.1.4 Loading the ICPCON Program Under Windows 95/98	43

C.1.5 Loading ICPCON Under SCO UNIX.....	43
C.1.6 Loading ICPCON Under LINUX.....	44
C.2 The ICPCON Program.....	44
C.2.1 Select Interface.....	45
C.2.2 Select Controller.....	46
C.2.3 The two Menu Areas „Monitor“ and „Express/Advanced Setup“.....	46
C.3 The Menu Monitor.....	47
C.3.1 Menu Monitor: View Statistics.....	47
C.3.2 Menu Monitor: View Events.....	48
C.3.3 Menu Monitor: View Hard Disk Info.....	48
C.3.4 Menu Monitor: Save Information.....	49
C.4 The Menu Express/Advanced Setup.....	50
C.4.1 Menu Express Setup: Configure Host Drives.....	50
C.4.2 Menu Express Setup: Repair Array Drives.....	54
C.4.3 Menu Advanced Setup: Configure Controller.....	56
C.4.4 Menu Advanced Setup: Configure Physical Devices.....	59
C.4.5 Menu Advanced Setup: Configure Logical Drives.....	65
C.4.6 Menu Advanced Setup: Configure Array Drives.....	66
C.4.7 Menu Advanced Setup: Configure Host Drives.....	71
D. ICP RAID NAVIGATOR.....	74
D.1 Introduction.....	74
D.2 The ICP RAID Navigator "Controls".....	75
D.2.1 The Toolbar.....	75
D.2.2 The Status Bar.....	75
D.2.3 "Window" Menu Commands.....	75
D.2.4 "Help" Menu Commands.....	76
D.2.5 "File" Menu Commands.....	76
D.2.6 "View" Menu Commands.....	76
D.2.7 The "Chart" Menu.....	77
D.2.8 The "Configuration" Menu Commands.....	77
D.3 Select Controller.....	78
D.4 Physical Configuration Window.....	78
D.4.1 Controllers.....	79
D.4.2 I/O Processors.....	81
D.4.3 Direct Access Devices.....	82
D.4.4 Non direct access devices (raw devices).....	86
D.5 Logical Configuration Window.....	87
D.5.1 The Host Drive Information Window.....	90

D.5.2	The Array Drive Information Window	90
D.5.3	The Logical Drive Information Window.....	92
D.5.4	Change the name of a Drive	93
D.5.5	Remove a Host Drive	93
D.5.6	Create a new Host Drive.....	93
D.5.7	Parity Verify.....	94
D.5.8	Parity Recalculate.....	94
D.5.9	Progress Information.....	95
D.5.10	Expansion of an Array	95
D.5.11	Add a Hot Fix Drive.....	96
D.5.12	Remove a Hot Fix Drive.....	97
D.5.13	Hot Fix Pool Access.....	97
D.5.14	Add a RAID 1 Component (Mirror a Drive).....	97
D.5.15	Remove a RAID 1 Component (Remove a Mirror Drive).....	98
D.5.16	Replace a Logical Drive	98
D.5.17	The Different States of an Array Drive	98
D.6	The Statistics Window	101
D.7	The Controller Events Window.....	102
D.8	ICP RAID Navigator Help	103
D.9	ICP Service and ICP Mail	104
E.	APPENDIX	110
E.1	Subsystem Enclosures Manufacturers	110
E.2	Index.....	111

Chapter A

General

Information

A. Introduction

In today's information technology world, High Availability Systems are gaining more and more importance. The demands on these systems are not only the continuous availability of all important data, i.e., "365 days * 24 hours", but also efficient resource sharing of the relatively expensive components. The secret behind the popular word "Clustering" is nothing more than the intelligent combination of several computer systems with the purpose to simultaneously increase the security and computing power of the whole system. This "Cluster" appears to the user just like one single system with extraordinary characteristics. A special Clustering software executes the controlling and management of the Cluster and its members.

There are several strategies to design Clustering systems. In a first model, the various members (also called nodes) could only be connected with each other through the normal User LAN. In another step, the nodes could additionally share an I/O bus (e.g., SCSI, Fibre Channel). To avoid "Single Points of Failure" (these are vital components which cause a general failure if they fail) in such systems, host based RAID controllers are ideally suitable for cluster systems. The choice of host based RAID controllers for the nodes of a cluster system implements the appropriate redundancy and security in the mass storage part of the system right from the beginning. In contrast to this, external RAID controllers carry the risk of becoming a single point of failure. If they fail the whole cluster system goes down. Apart from that, these kinds of solutions are normally pretty inhomogeneous: The host adapters come from vendor A, the external RAID controller from vendor B. Who guarantees the flawless and interference-free operation of both components? Who supports you when the system causes problems?

Since 1990 ICP vortex has been intensively engaged in the research and development of host based RAID products for highest performance and security requirements. Not only because of our products' outstanding performance, our expertise and continuity in development, but also due to our excellent customer support, ICP Controllers are accepted and known as top leading-edge products all over the world. All ICP Controller products are pure-bred **hardware solutions**. All functionality required for the very complex tasks is hardware-implemented on the controller. Thus, RAID is fully independent of the computer system (the host) and the operating system. Software RAID solutions are normally not usable for cluster systems. Thanks to the wide operating system support and easy-to-use installation and maintenance utilities, setting up and using high performance and fault-tolerant mass-storage subsystems for almost every purpose is child's play.

For a successful installation of a Clustering System it is very important that the person who carries out the installation has detailed knowledge and experience in the areas Microsoft® Windows NT/2000 Server, Microsoft® Cluster Server, Novell/Orion, Linux/Watchdog and ICP Disk Array Controllers.

This Addendum contains important information, particular regarding the operation of ICP RAID Controllers in a Clustering System.

Information on the Operating System, the functioning of the Cluster and other system details may be found in the corresponding system manuals.

For information on the installation of the ICP Disk Array Controller and the configuration of Disk Arrays, please refer to the "User's Manual" or "Quick Installation Guide" of your ICP Controller. These documents may also be found on the ICP System CDROM.

We would like to thank you for purchasing an ICP Controller.

ICP - Intelligent Computer Peripherals®



A.1 ICP Controllers in a Cluster Configuration

Due to the required cable lengths (several meters) for the shared I/O channel(s) and the demands on the transfer performance and quality, only Wide/Ultra2 SCSI, Ultra 160 SCSI or Fibre Channel controllers are suitable for Clustering. Using such controllers allows for sufficient distances between the two servers and the RAID enclosure.

All ICP Controllers with **Clustering functionality** have a **6** after the first number of the product name. ICP Controllers currently available with Clustering functionality are:

Wide/Ultra2 SCSI, 32-Bit PCI: GDT6618RD, GDT6628RD, GDT6638RD, GDT6658RD
Wide/Ultra2 SCSI, 64-Bit PCI: GDT7618RN, GDT7628RN, GDT7638RN, GDT7658RN
Ultra 160 SCSI, 64-Bit PCI: GDT7623RN, GDT7643RN, GDT7663RN
Fibre Channel, 64-Bit PCI: GDT7619RN, GDT7629RN

By using the **Cluster Module** (Part No.: 8780), every 65.. and 75.. ICP Wide/Ultra2 SCSI, Ultra 160 SCSI or Fibre Channel Controller can be upgraded to an ICP Controller with Clustering functionality.

ICP Part Number	ICP Controller Name	Number and Type of I/O Channels	Supported RAID Levels
7867	GDT6618RD	1 * WU2 / 1 * NU	0/1/4/5/10
7877	GDT6628RD	2 * WU2 / 1 * NU	0/1/4/5/10
7887	GDT6638RD	3 * WU2 / 1 * NU	0/1/4/5/10
7897	GDT6658RD	5 * WU2 / 1 * NU	0/1/4/5/10
7868	GDT7618RN	1 * WU2 / 1 * WU	0/1/4/5/10
7878	GDT7628RN	2 * WU2 / 1 * WU	0/1/4/5/10
7888	GDT7638RN	3 * WU2 / 1 * WU	0/1/4/5/10
9607	GDT7623RN	2 * U160	0/1/4/5/10
9608	GDT7643RN	4 * U160	0/1/4/5/10
9609	GDT7663RN	6 * U160	0/1/4/5/10
7987	GDT7619RN	1 * FC / 1 * WU	0/1/4/5/10
7997	GDT7629RN	2 * FC / 1 * WU	0/1/4/5/10

WU2 = Wide/Ultra2 SCSI / U160 = Ultra160 SCSI / FC = Fibre Channel Arbitrated Loop / WU = Wide/Ultra / NU = Narrow/Ultra

For the successful installation and setup of Clustering Systems not only the computer systems, the RAID Controllers and mass storage enclosures are important, but also the accessories, such as cables, terminators, Fibre Channel Hubs and SCSI Bus Repeaters, Isolators or Converters. ICP vortex offers a variety of genuine high quality accessories, which are highly recommended for Clustering Systems.

For external mass storage enclosures we have tested various products. "Fortra" Series enclosures of the company JMR are pretty often used. A list of vendors can be found in the appendix of this addendum.

A.2 Copyrights, Patents

Parts of the ICP GDT Rx Series controllers are protected under international copyright laws and agreements. No part of the product or the manual, or parts of the manual may be reproduced in any form, physical, electronic, photographic, or otherwise, without the expressed written consent of ICP vortex Computersysteme GmbH. For this product a patent is registered at the Deutsches Patentamt in Munich with the official reference no. 4121974.

All special names and trademarks of manufacturers quoted in this manual are protected by copyright.

ICP - Intelligent Computer Peripherals ® and **RAIDYNE** ®, are registered trademarks of ICP vortex Computersysteme GmbH.

Europe: ICP vortex Computersysteme GmbH / Konrad-Zuse-Str. 9 / 74172 Neckarsulm - Germany / Tel.: +49-(0)-7132-9620-0 / Fax: +49-(0)-7132-9620-200 / E-Mail: support@vortex.de / WWW: <http://www.icp-vortex.com>

United States of America: ICP vortex Corporation / 4001 E. Broadway, B-20 / 85040 Phoenix, Arizona / Phone: 602-414-0414 / Fax: 602-414-0444 / E-Mail: support@icp-vortex.com / WWW: <http://www.icp-vortex.com>

ICP vortex is member of the following associations and initiatives:



A.3 Software License Agreement

Please read this Software License Agreement before opening the CD/disk packaging and before starting to use the programs. Each loading of a program covered by this license agreement, each transmission within any existing network to another computer, as well as each copy on a mass storage system, regardless of what kind (floppy disk, hard disk, CD, MO, etc.), represents a duplication of the program according to copyright regulations. Duplication is permitted only with the authorization of ICP vortex.

This authorization will be granted only on the condition that the Software License Agreement stated hereafter is observed.

By opening the CD/disk packaging you expressly acknowledge the Software License Agreement of ICP vortex.

1. You are authorized to use the software contained on the enclosed disks, CDROMs and EPROMs/Flash-RAMs on a single computer system only. The restriction to this one computer system also applies if the disk packaging contains a double set of software, for example one set of 3.5" floppy disks and a CDROM. It is further valid if the package contains several versions of software adapted to different operating systems. Multi-utilization of the software is only permitted if a multi-user license has been purchased. The number of further computer systems authorized for usage under a multi-user license is evident from and limited by that license.
2. It is permitted to produce one single copy disk of the software for back-up purposes only. Furthermore, it is permitted to copy the software onto the hard disk of one single computer. It is not permitted to duplicate the contents of the EPROMs and/or Flash-RAMs on the ICP Controller.
3. The permanent conferring (by sale or donation) of the software is permitted. The new proprietor must be registered with ICP vortex and must assume all rights and obligations resulting from this Software license agreement. Each and any other kind of transfer, in particular leasing, is not permitted. Copies made by the first user for security reasons must be destroyed upon transfer.
4. It is not allowed to change the software in its functions or its appearance (especially trade mark, firm name and copyright reference) or to edit it in any other way. Neither is it permitted to de-compile or disassemble the software.
5. The enclosed software has been carefully copied on floppy disks and/or CDROM(s). However, if the floppy disks and/or CDROM(s) should prove to be faulty, ICP vortex will exchange them within 4 weeks of the date of purchase.
6. ICP vortex makes no warranties, express or implied, including without limitation the implied warranties of merchantability, functionality and fitness for a particular purpose. Furthermore, ICP vortex is not liable to you for any consequential, incidental or indirect damage arising out of the use of this product.
7. This agreement is subject to the laws of the Federal Republic of Germany. Place of jurisdiction for both parties is the domicile of ICP vortex Computersysteme GmbH.

A.4 ICP vortex Accessories for SCSI Controllers

ICP vortex offers various SCSI accessories for the setup and installation of Cluster Systems based on SCSI technology. These accessories are not only usable for SCSI ICP RAID Controllers, but also for Controllers with SCSI interfaces from other vendors.

A.4.1 SCSI Wonder

SCSI Wonder is a very powerful and universal add-on for SCSI bus systems. It was developed with a special focus on Cluster Systems with parallel SCSI busses, but may also be used for totally different tasks (also in combination with SCSI RAID Controllers or SCSI Host Adapters of other vendors). SCSI Wonder allows a fully transparent connection between two SCSI bus segments. Both segments can be based on different transfer technologies (SE or LVD). SCSI Wonder does not require a SCSI ID on the SCSI bus and has no SCSI bus protocol functions. SCSI Wonder is for both, the initiator (normally the SCSI RAID Controller) and the targets (the hard disks, CDRoms, DLTs, DATs, etc.) „invisible“.

SCSI Wonder can take over the following tasks:

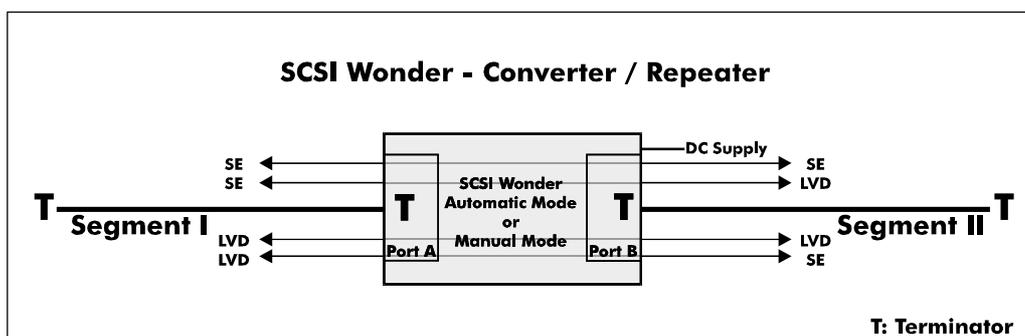
- **Extension of an existing SCSI bus segment (up to 3 SCSI Wonder can be used on one SCSI bus):**

Connection of two LVD SCSI bus segments (LVD <-> LVD)

Connection of two SE SCSI bus segments (SE <-> SE)

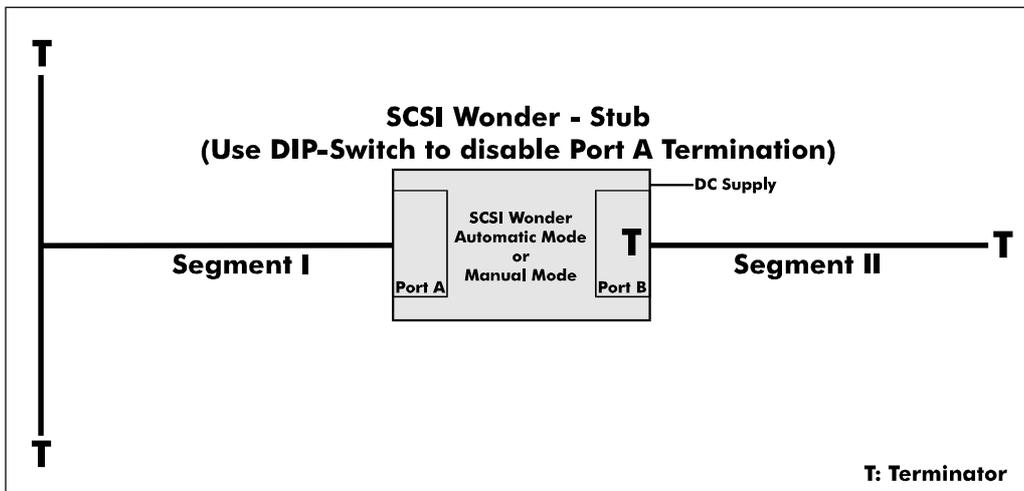
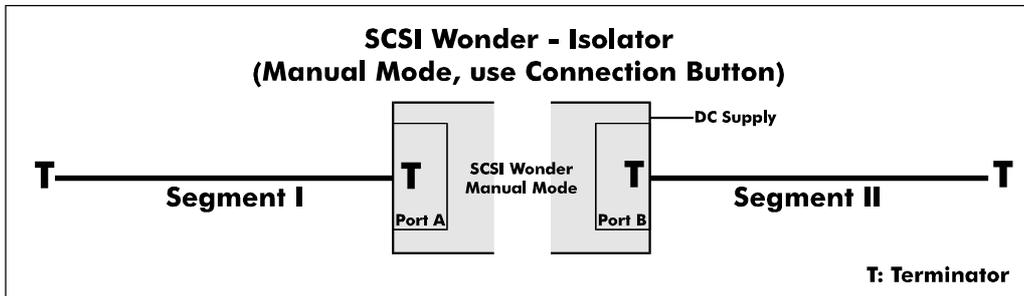
Conversion of a SE SCSI bus segment into an LVD SCSI bus segment and vice versa (SE <-> LVD; LVD <-> SE)

Installation of T-connections (stubs)



- **Isolation of two SCSI bus segments:**

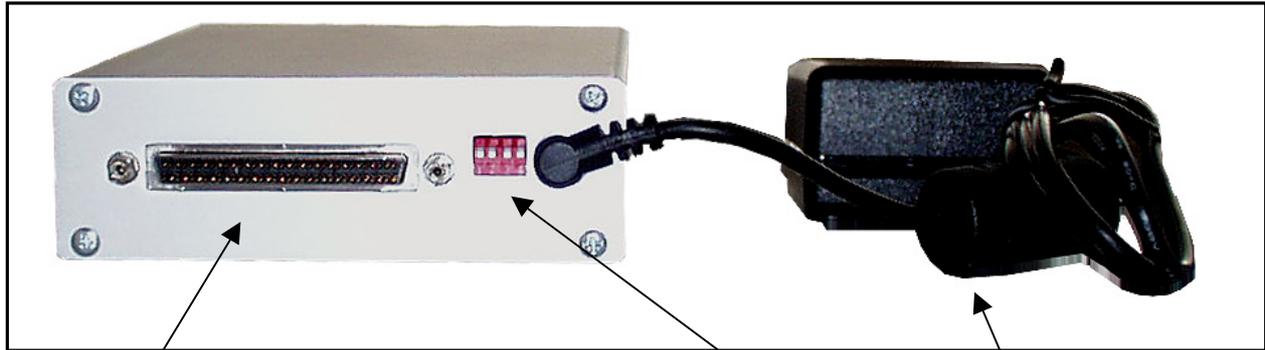
This brings a lot of security for Cluster Systems; SCSI bus segments can be separated or reconnected without any detrimental effects for the SCSI bus (i.e., the other vital Server node).



SCSI Wonder as an Isolator fulfills two primary tasks. First, it maintains a proper termination of the shared SCSI bus, even if one Node is removed. Second, it keeps the SCSI bus signal quality on all remaining segments on an extremely good and healthy level.

A.4.2 SCSI Wonder Overview

(Note: This description is also printed on the bottom of the SCSI Wonder enclosure.)



Port B - SE/LVD

Automatic Mode (Factory Setting)
All DIP-Switches down.



Manual Mode (use Connection Button)



Port A: Termination ON
Port B: Termination OFF



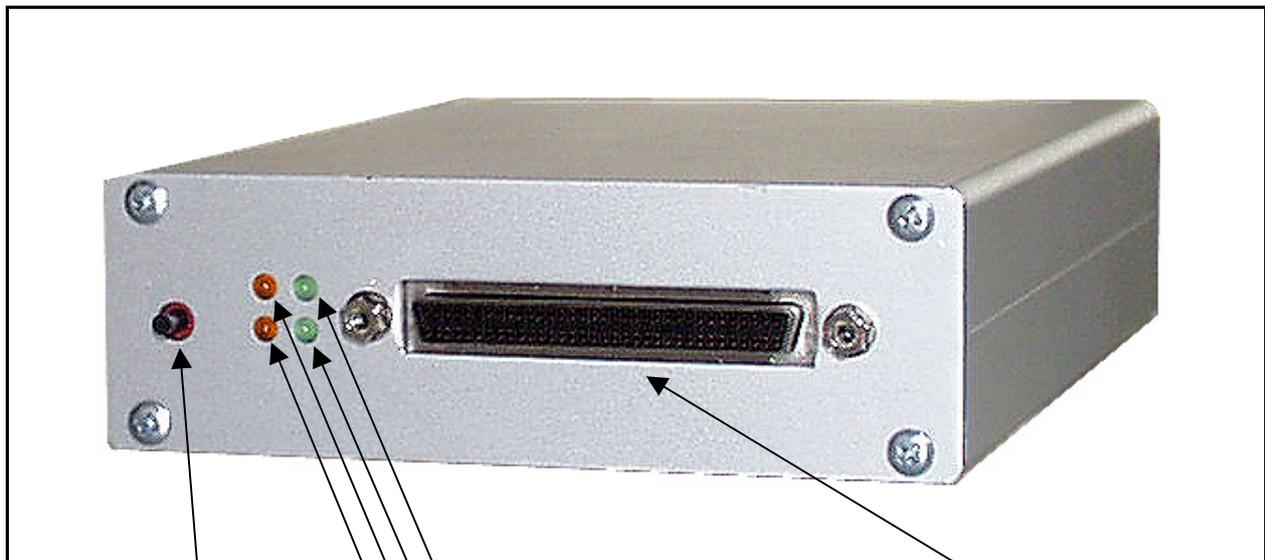
Port A: Termination OFF
Port B: Termination ON



Port A: Termination OFF
Port B: Termination OFF



ICP vortex Power Supply.
5V/4A DC.
Use only the original ICP vortex part !



Connection Button (only in Manual Mode available).
Press once: Ports A and B isolated.
Press once again: Ports A and B connected again.

Termpower Port B (ON=Termpower available)

Termpower Port A (ON=Termpower available)

SCSI Bus Activity

Isolation Mode (OFF=Ports A and B connected)



Port A - SE/LVD

A.4.3 Cables, Terminators, Brackets

For the configuration of stable, high performance SCSI-Systems, we highly recommend you to use genuine ICP vortex accessories.

Order #	Part Name	Description	Application
8839	Wide/Ultra2 & Ultra 160 Flat Ribbon Cable	150 cm Wide/Ultra2 & Ultra 160 SCSI cable with seven 68 pin HD SCSI connectors (male).	Connection of up to 5 internal Wide/Ultra SCSI devices per SCSI channel and one SE/LVD terminator (Order No.: 8951). For LVDS operation.
8844	Wide/Ultra2 & Ultra 160 Round Cable	3 meters Wide/Ultra2 & Ultra 160 Round Cable with stackable 68 pin VHD SCSI connector (male) and 68 pin HD SCSI connector (male). Both with thumb screws.	Connection of an external Wide/Ultra2 & Ultra 160 SCSI subsystem. For LVDS operation.
8877	Wide/Ultra2 & Ultra 160 Round Cable	1 meter Wide/Ultra2 & Ultra 160 Round Cable with stackable 68 pin VHD SCSI connector (male) and 68 pin HD SCSI connector (male). Both with thumb screws.	Connection of an external Wide/Ultra2 & Ultra 160 SCSI subsystem. For LVDS operation.
8845	Wide/Ultra2 & Ultra 160 SCSI Bracket	Bracket with two internal 68 pin HD SCSI connector (female) and two external 68 pin VHD SCSI connectors (female).	Connection of external Wide/Ultra2 & Ultra 160 SCSI subsystems. For LVDS operation.
8849	Wide/Ultra2 & Ultra 160 SCSI Bracket	Bracket with one internal 68 pin HD SCSI connector (female) and one external 68 pin VHD SCSI connector (female).	Connection of an external Wide/Ultra2 & Ultra 160 SCSI subsystem. For LVDS operation.
8850	Wide/Ultra2 & Ultra 160 SCSI Bracket	Bracket with one internal 68 pin HD SCSI connector (female) and one external 68 pin HD SCSI connector (female).	Connection of an external Wide/Ultra2 & Ultra 160 SCSI subsystem. For LVDS operation.
8880	Wide/Ultra2 & Ultra 160 Round Cable	1 meter Wide/Ultra2 & Ultra 160 Round Cable with 68 pin HD SCSI connectors (male). Both with thumb screws.	Connection of an external Wide/Ultra2 & Ultra 160 SCSI subsystem. For LVDS operation.
8951	LVDS/SE Terminator	Terminator for LVDS or SE buses with 68 pin HD SCSI connector (female). LEDs indicate, whether the bus/terminator is operating in SE or LVDS mode.	Termination of SE or LVDS 68 pin cables. For LVDS operation.
8952	LVDS/SE Terminator for SE/LVDS Hard Disks with SCA Connector.	Terminator for LVDS or SE buses with SCA connector for the hard disk and 68 pin HD SCSI connector (female). Connector for DC. Jumpers for SCSI-ID, Termination, Termpower, etc.	This product converts Wide/Ultra2 & Ultra 160 or Wide/Ultra hard disks with SCA connector into "normal" hard disks with termination (SE and LVDS), DC connector, Jumpers for SCSI-ID, etc.. For LVDS operation.
8953	External LVDS/SE Terminator	Stackable external Terminator for SE or LVD busses with 68 pin VHD SCSI connector (male). Thumb-Screws.	For LVDS operation.

4.4.4 SCSI Wonder to set up Cluster-Systems

For the configuration of stable, high performance Cluster-Systems, we highly recommend you to use genuine ICP vortex accessories.

Equipment list per shared SCSI channel (from the controllers to the shared disk enclosure).

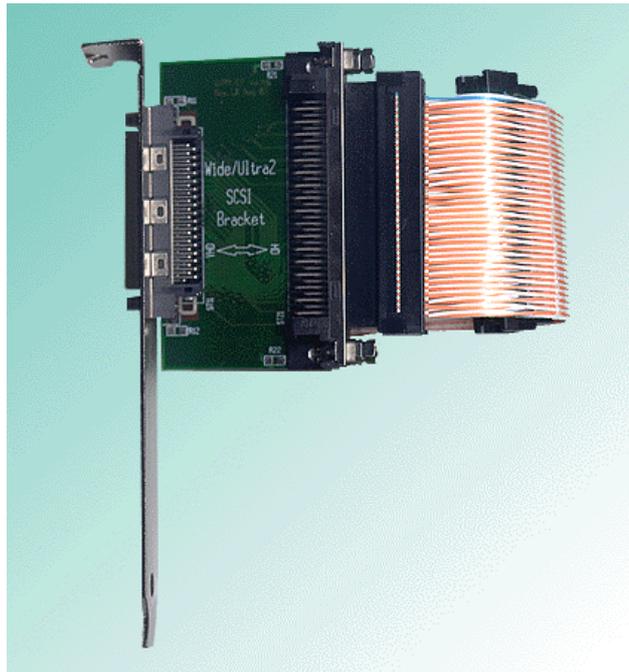
ICP vortex Part-Number	Description	Number
8954	SCSI Wonder incl. Power supply	2
8844 or 8877	1 Meter / 3 Meter from the controller to SCSI Wonder	2
8844, 8877 or 8880	1 Meter / 3 Meter from SCSI Wonder to the shared disk enclosure	2



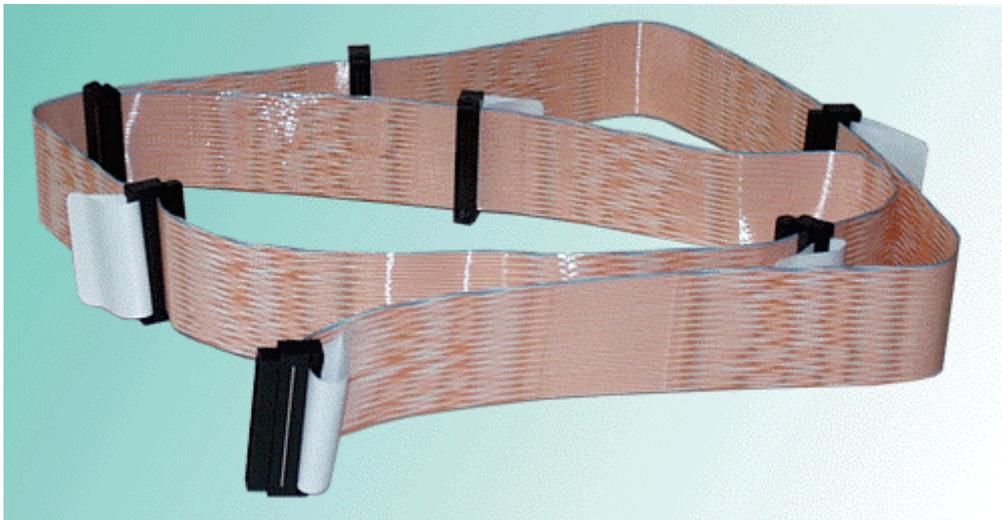
(8844, 8877. 1 Meters, 3 Meters external Wide/Ultra2 SCSI cables)

Additional equipment list per shared SCSI channel for the cabling inside the shared disk enclosure (if necessary).

ICP vortex Part-Number	Description	Number
8849	Wide/Ultra2 SCSI Bracket VHD<->HD IN-and OUT connectors for the subsystem	2
8839	Internal Wide/Ultra2 flat ribbon cable for 5 hard disks	1

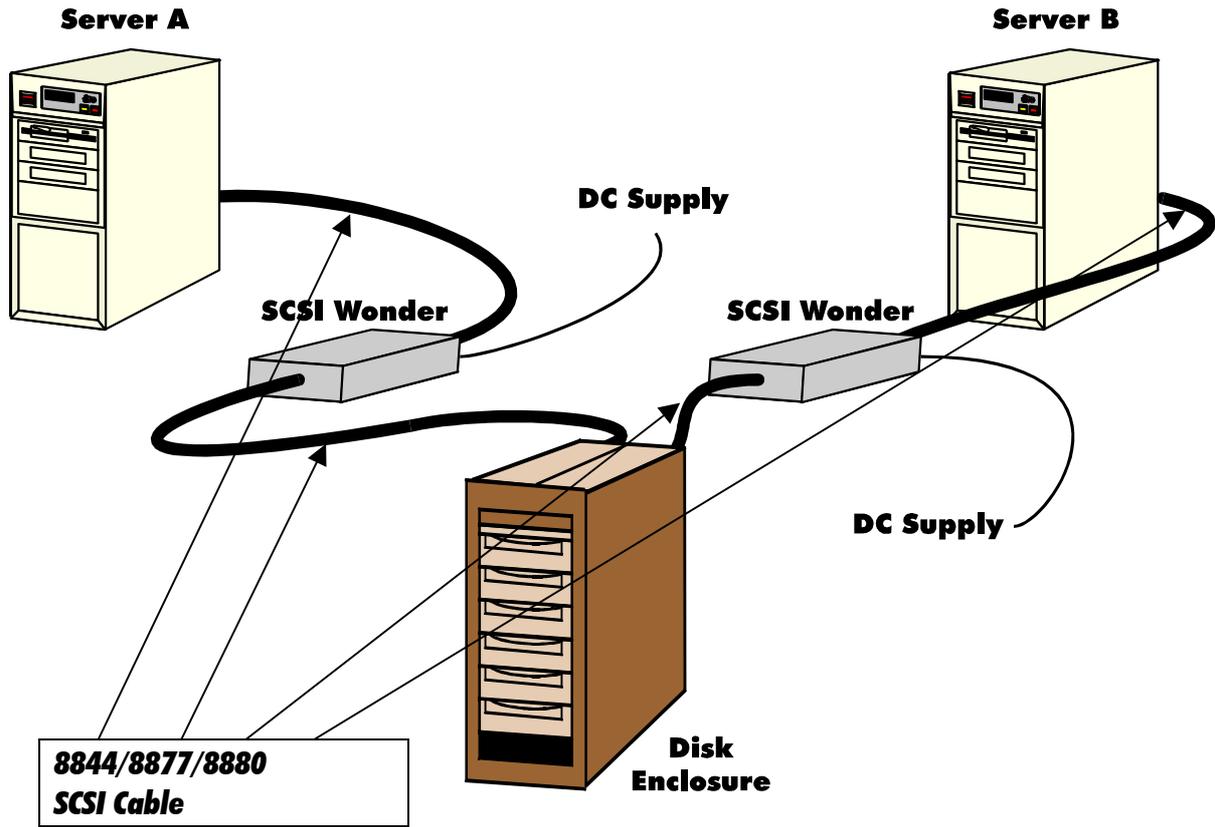


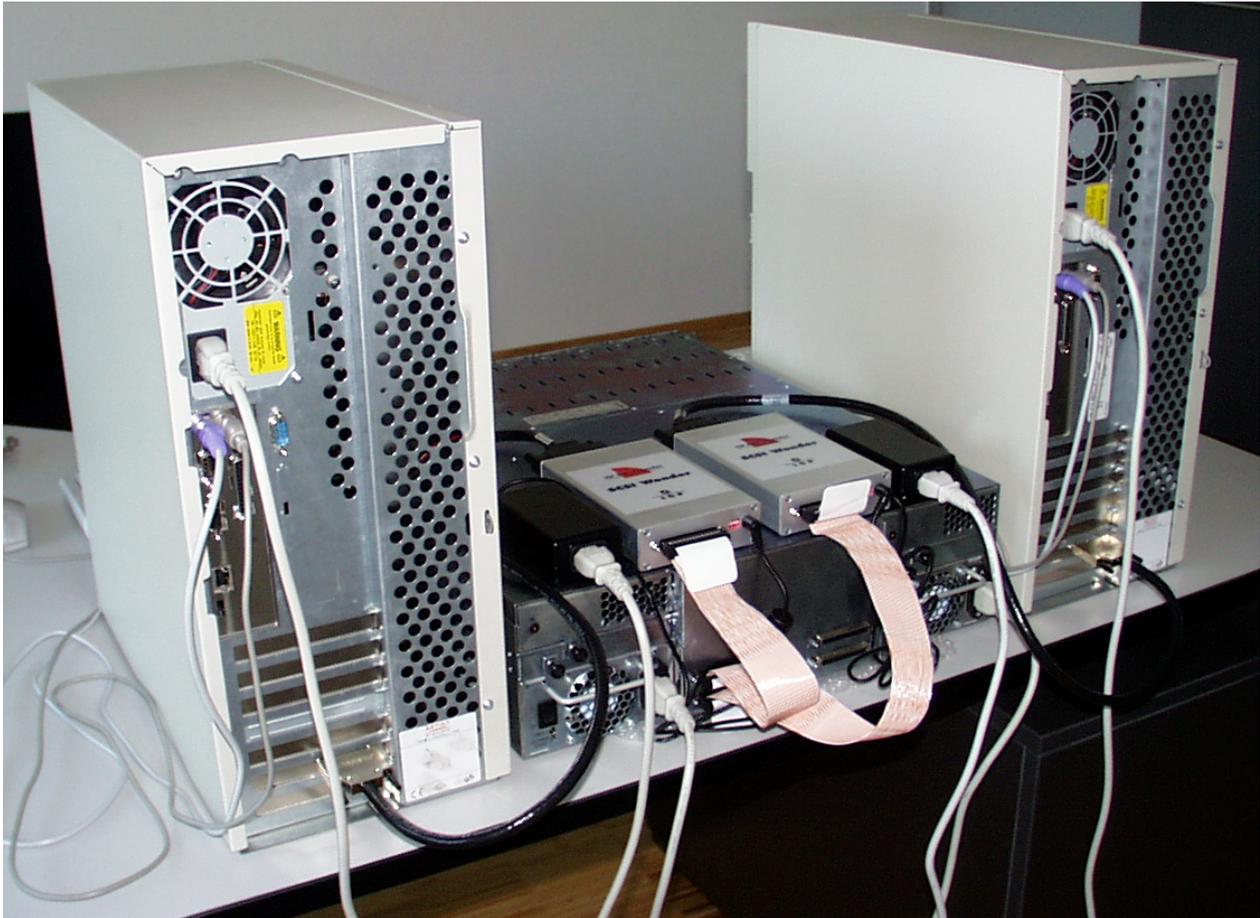
(8849. Wide/Ultra2, Ultra 160 SCSI Bracket)



(8839. Wide/Ultra2, Ultra 160 SCSI flat ribbon cable)

Typical Cluster Configuration:





A.5 ICP vortex Accessories for Fibre Channel Controllers

A.5.1 ICP Fibre Hub 1016

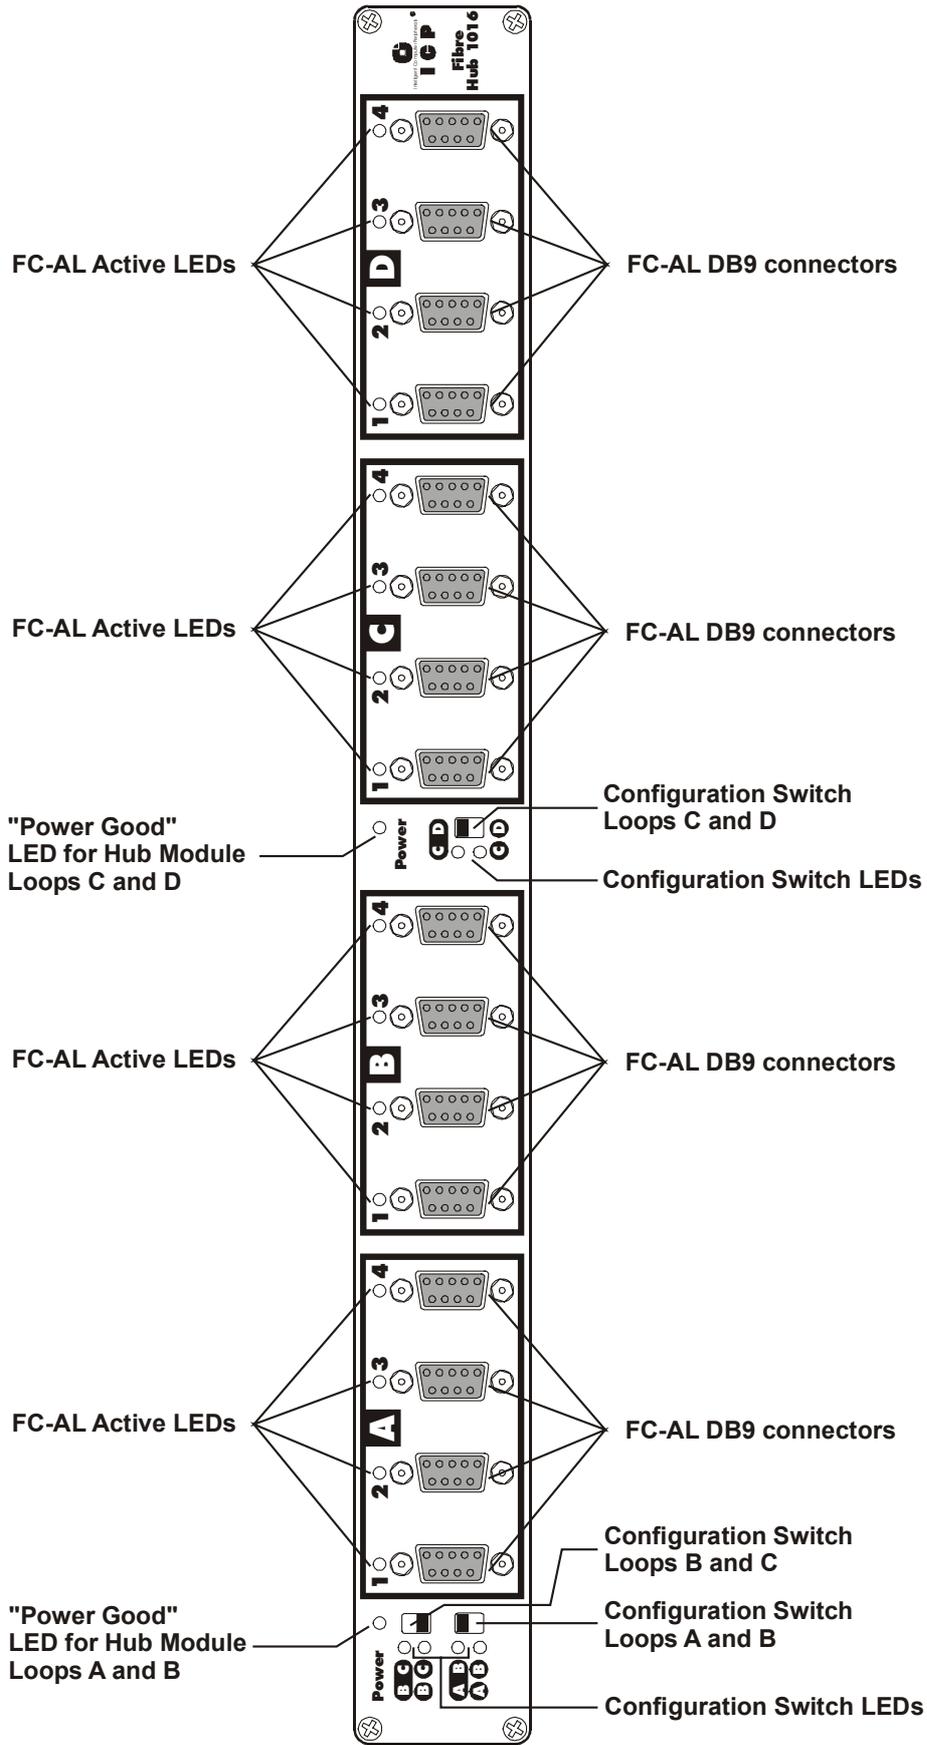
Different to other solutions, the ICP Fibre Hub 1016 is designed not only for typical configurations which require a hub, but also for mission critical SAN and Clustering applications.

The 19" enclosure hosts **two fully independent Hub Modules** and **two fully independent power supplies**. Each Hub Module contains two loops with 4 ports per loop. With ICP Fibre Hub 1016 in a Dual Loop setup not only the loops are redundant, but also the hubs and power supplies. The ICP Fibre Hub 1016 is able to connect up to 16 FC-AL devices in a physical star, while keeping the devices in a logical ring configuration.

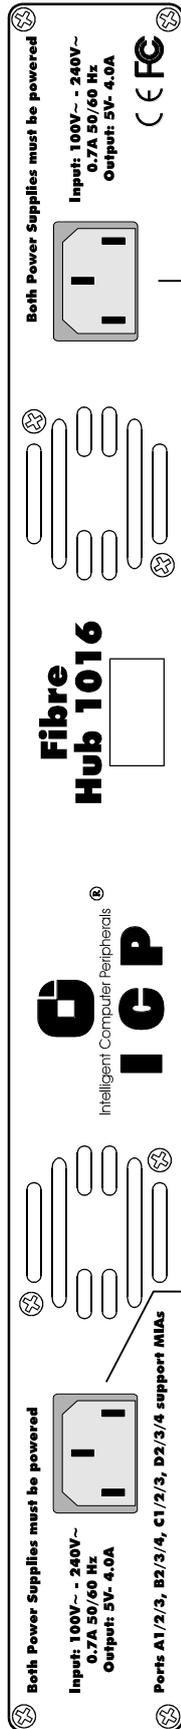
Connections to the hub use DB-9 connectors. Non-operating devices are automatically bypassed and newly communicating devices are automatically entered onto the loop. Full Duplex operation per port with 1.0625 Gbit/sec. Repeater function per Port. The ICP Fibre Hub 1016 can be symmetrically divided into 2, 3 or 4 separate loops with 4 ports in each loop by using configuration switches.

12 of the 16 ports support the use of external Media Interface Adapters (MIAs)

ICP Fibre Hub 1016 – Front View



ICP Fibre Hub 1016 – Rear View

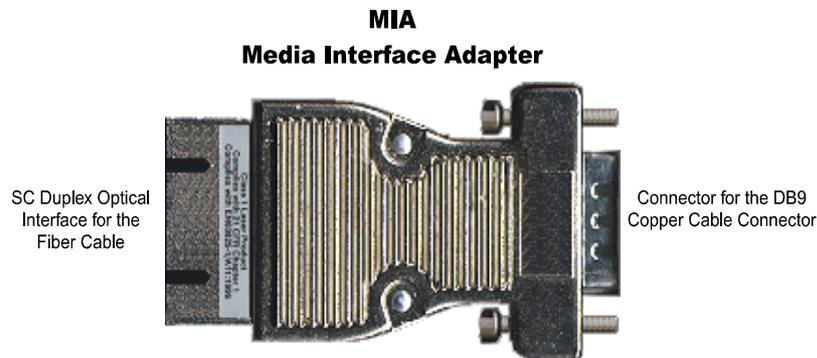


AC Input: 100V~ - 240V~
0.7A 50/60 Hz
Output: 5V- 4.0A

AC Input: 100V~ - 240V~
0.7A 50/60 Hz
Output: 5V- 4.0A

A.5.2 ICP vortex Fibre Channel Cables and MIAs

MIA adapters transform the electrical signals used by copper cables into optical signals transmitted by optical fibers. One end of the adapter has an FC DB 9 male connector and the other end has an SC duplex female connector for the optical fiber. The adapter uses a laser diode to transform electrical signals into light signals and an opti-



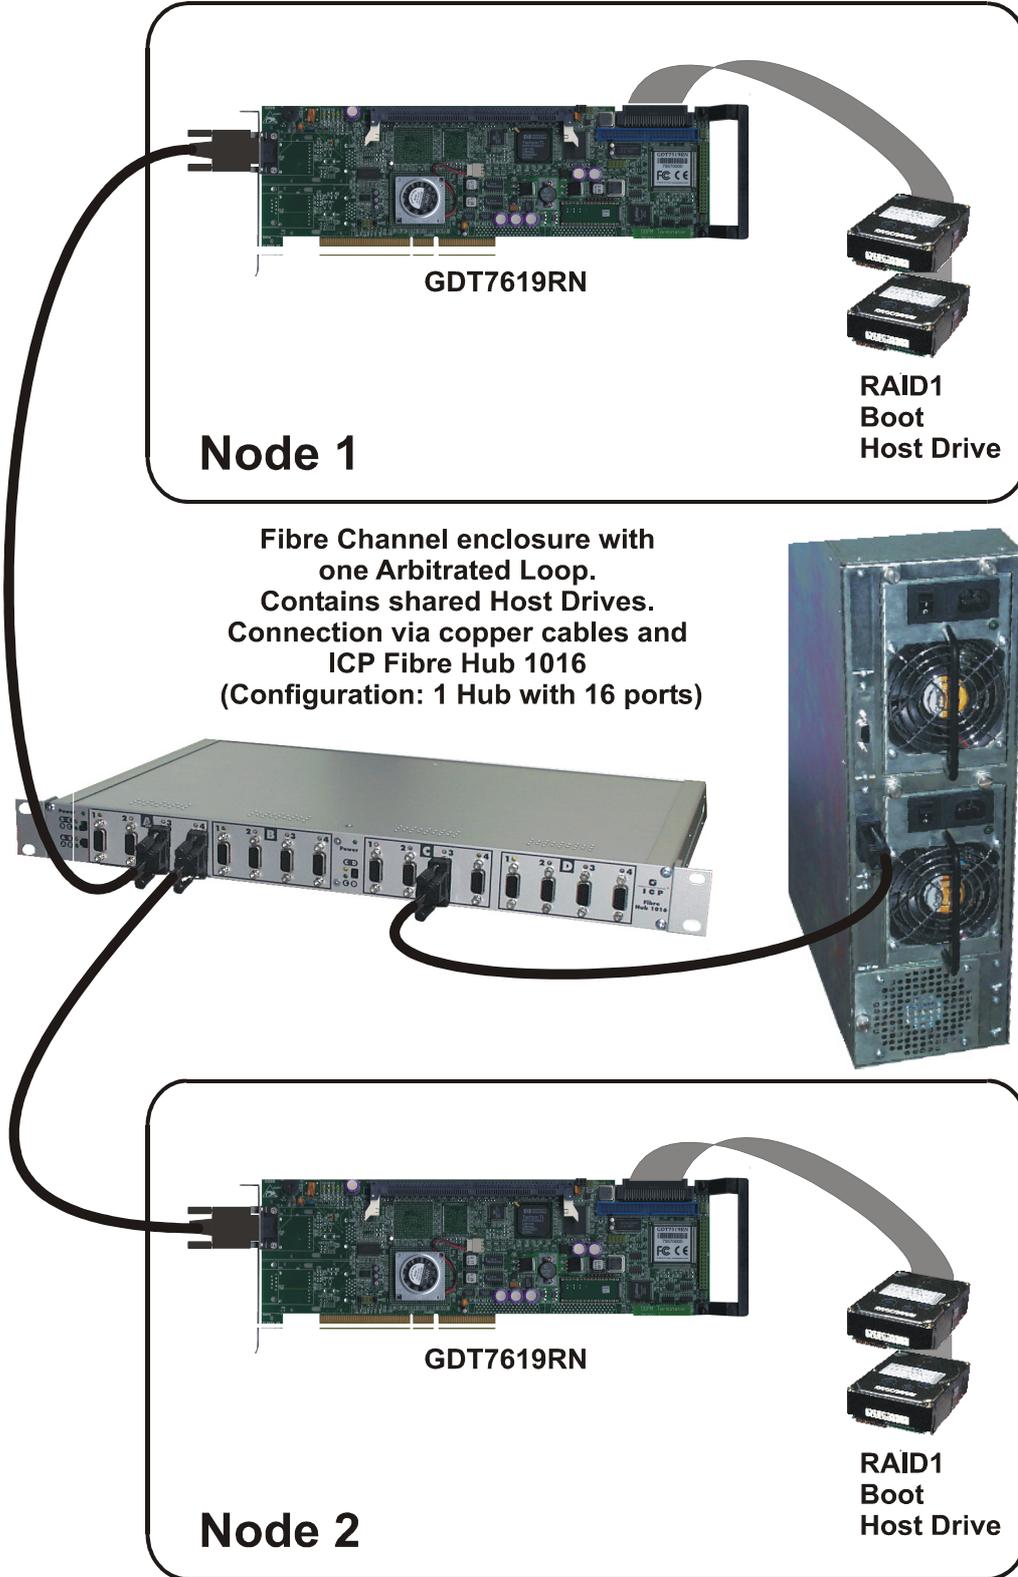
cal sensor to perform the reverse function. 12 of the 16 ports of the ICP Fibre Hub 1016 support MIAs.

3 Meters Fibre Channel Copper Cable.

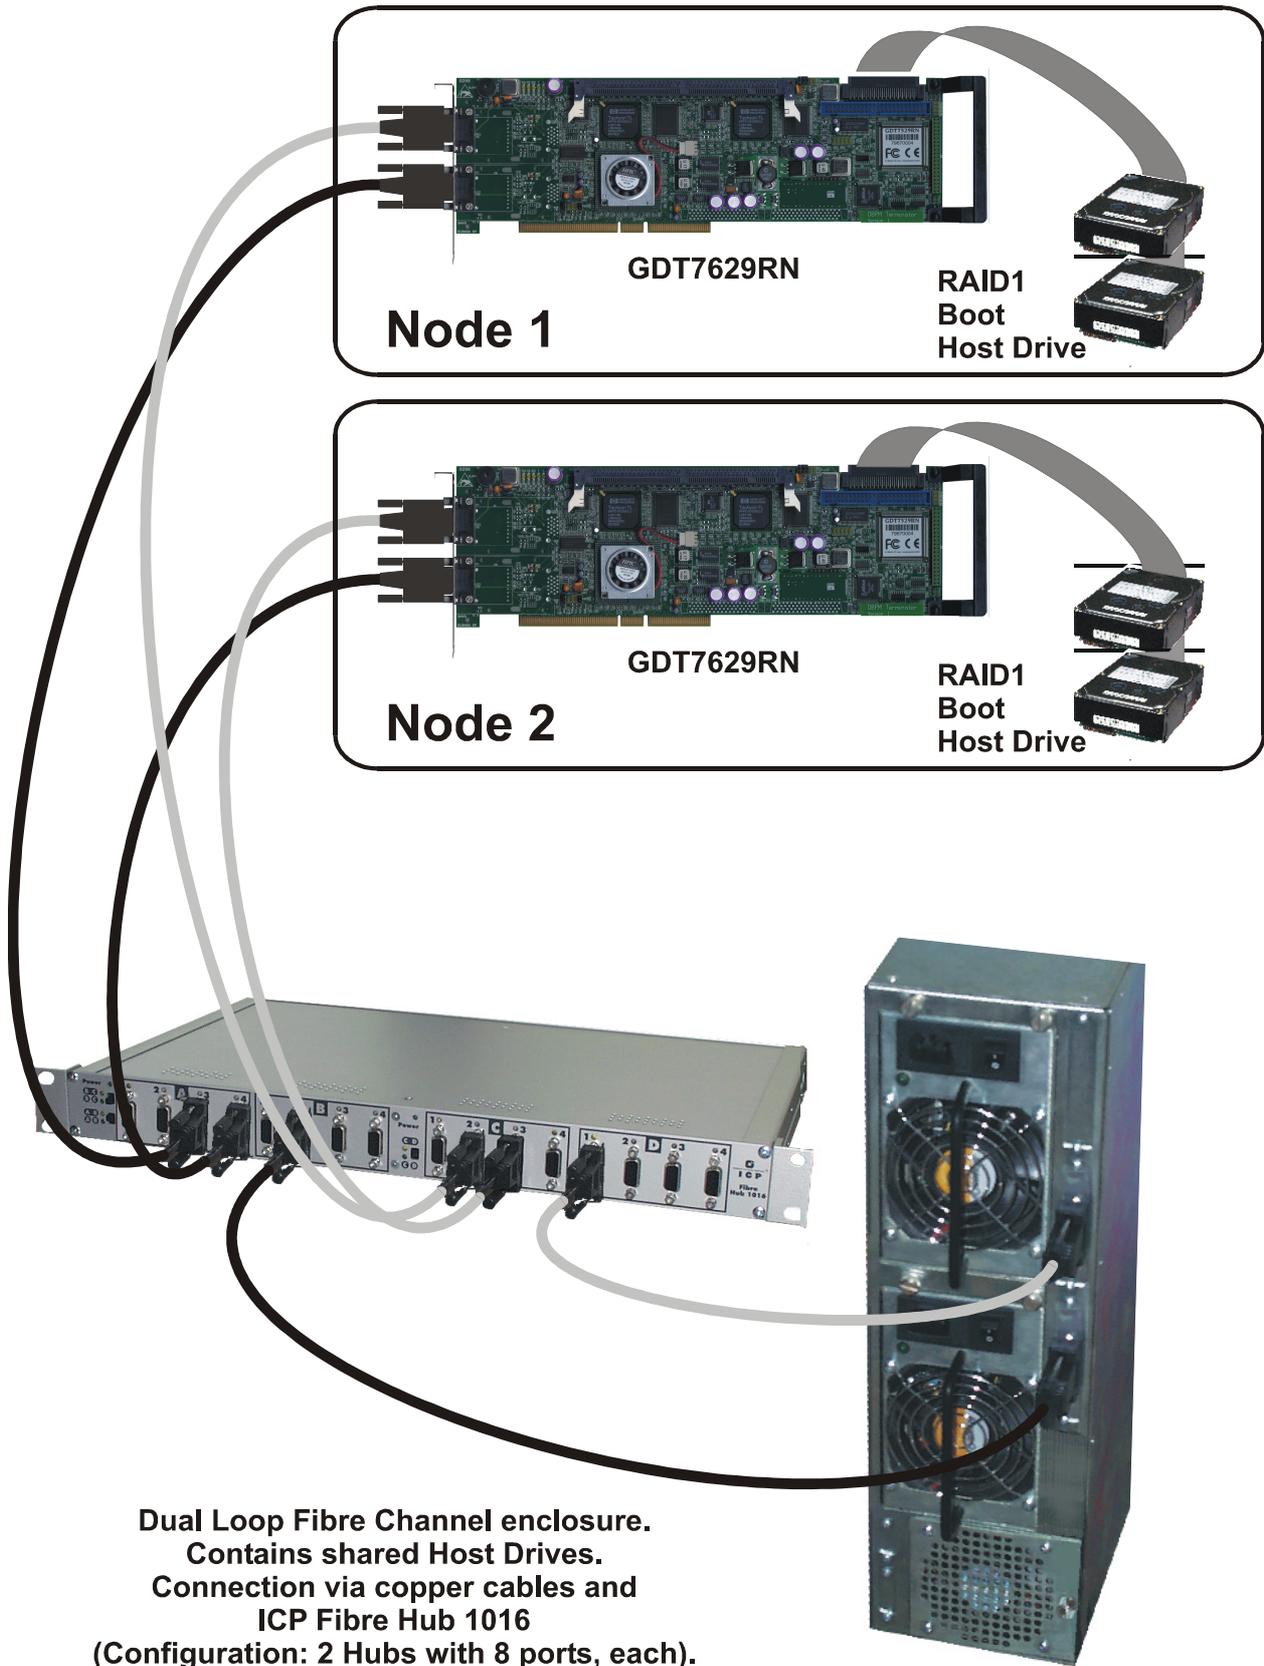


A.5.3 Sample Configurations

Clustering with one FC-AL; ICP Fibre Hub 1016 configured as one Hub with 16 ports



Clustering with Dual Loop FC-AL; ICP Fibre Hub 1016 configured as two independent Hubs with 8 ports, each. Note: In this configuration Fibre Hub 1016 is split into two fully independent Hub Modules, each with a separate power supply.



Chapter B

Installation Notes

B. Installation Notes

B.1 Hardware Installation

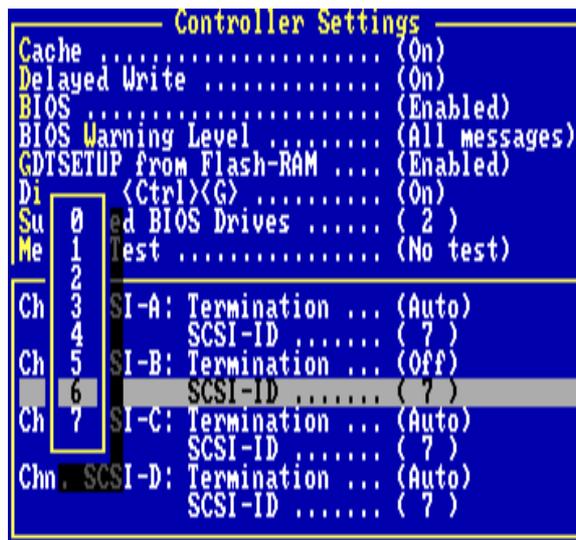
Wide/Ultra2 SCSI or Ultra 160 SCSI channels which are used as a shared I/O bus require per shared channel 2 SCSI Wonders to make sure that the shared channel continues to work even if a server is switched off, or if the bus is separated from the ICP Controller in one of the servers.

For the configuration of Cluster Systems with ICP Fibre Channel Controllers, it is recommended to use the ICP Fibre Hub 1016, which ensures that the loop is always closed even if one Controller fails, or if the fibre channel cable is removed. ICP Controllers with two independent loops allow the set up of a dual loop configuration. In such a configuration, there is not only redundancy at server level but also at I/O channel level: One loop can fail and the system continues to run on the other loop. ICP Fibre Hub 1016 was specifically designed for such dual loop configurations (it consists of two independent Hub Modules and two independent power supplies).

B.2 Using Wide/Ultra2 and Ultra 160 SCSI ICP Controllers

Both ICP Controllers for the two nodes should be identical (same model, same firmware level).

Before connecting both controllers to the shared I/O cable, one of the IOPs (Input/Output Processor) on one of the controllers must be set to another ID e.g., ID6). This has to be done for all I/O channels which are used as Clustering channels. The easiest way to change the IOP IDs is to use ICPCON (ICP RAID Console) loaded with <CTRL>+<G> from the controller's ROM.



B.3 Using Fibre Channel ICP Controllers

Both ICP Controllers for the two nodes should be identical (same model, same firmware level).

We recommend to use for the shared Fibre Channel loops the ICP Fibre Hub 1016. This Hub can be split into two independent Hubs with own power supplies and 8 ports per hub. This configuration should be used if a Dual Loop Configuration is installed.

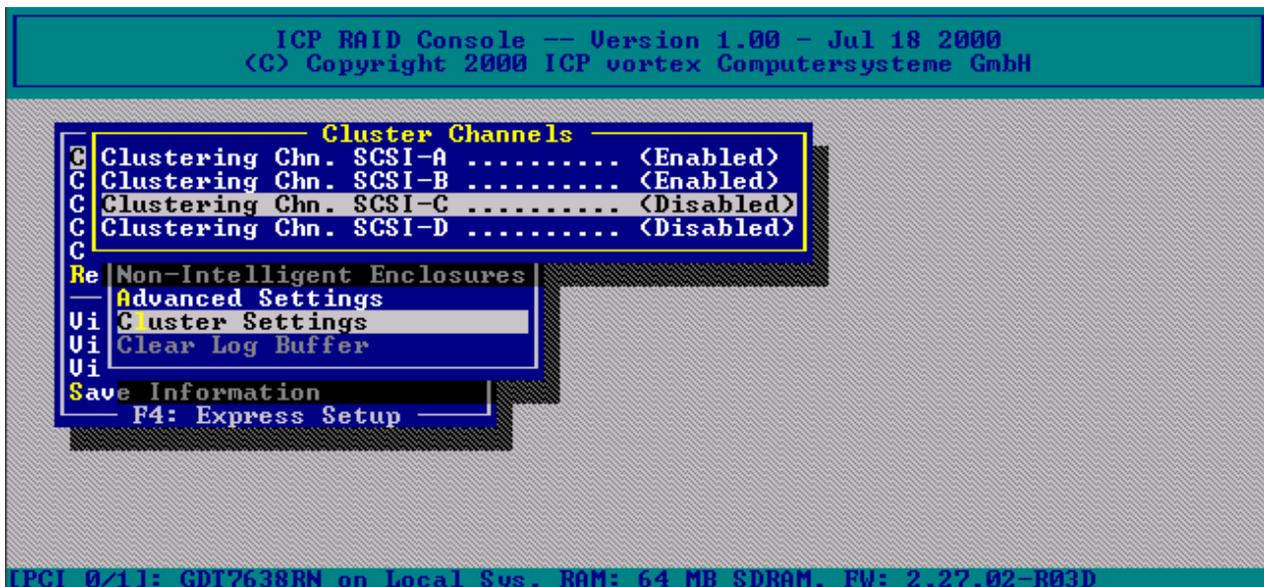
When choosing an external Fibre Channel enclosure make sure that it is certified for Dual Loop operation and has DB9 connectors which support MIAs (Media Interface Adapter) for configuring optical cabling. The same also applies to the DB9 connectors of the Fibre Channel Hubs .

The appendix lists various enclosure manufacturers whose products have been successfully tested in our labs.

Unlike the IDs of the Wide/Ultra2 or Ultra 160 SCSI IOPs, the IDs of the Fibre Channel IOPs of the ICP Controllers are automatically set to different numbers when operated on one loop.

B.4 Releasing the Cluster Channels

Those channels, which are used as shared channels for Clustering, have to be released and marked as Clustering Channels within the ICP RAID Console (press <CTRL>+<G> at system boot level when the ICP BIOS is displayed.). This applies to all ICP Controllers participating in the shared channel(s).

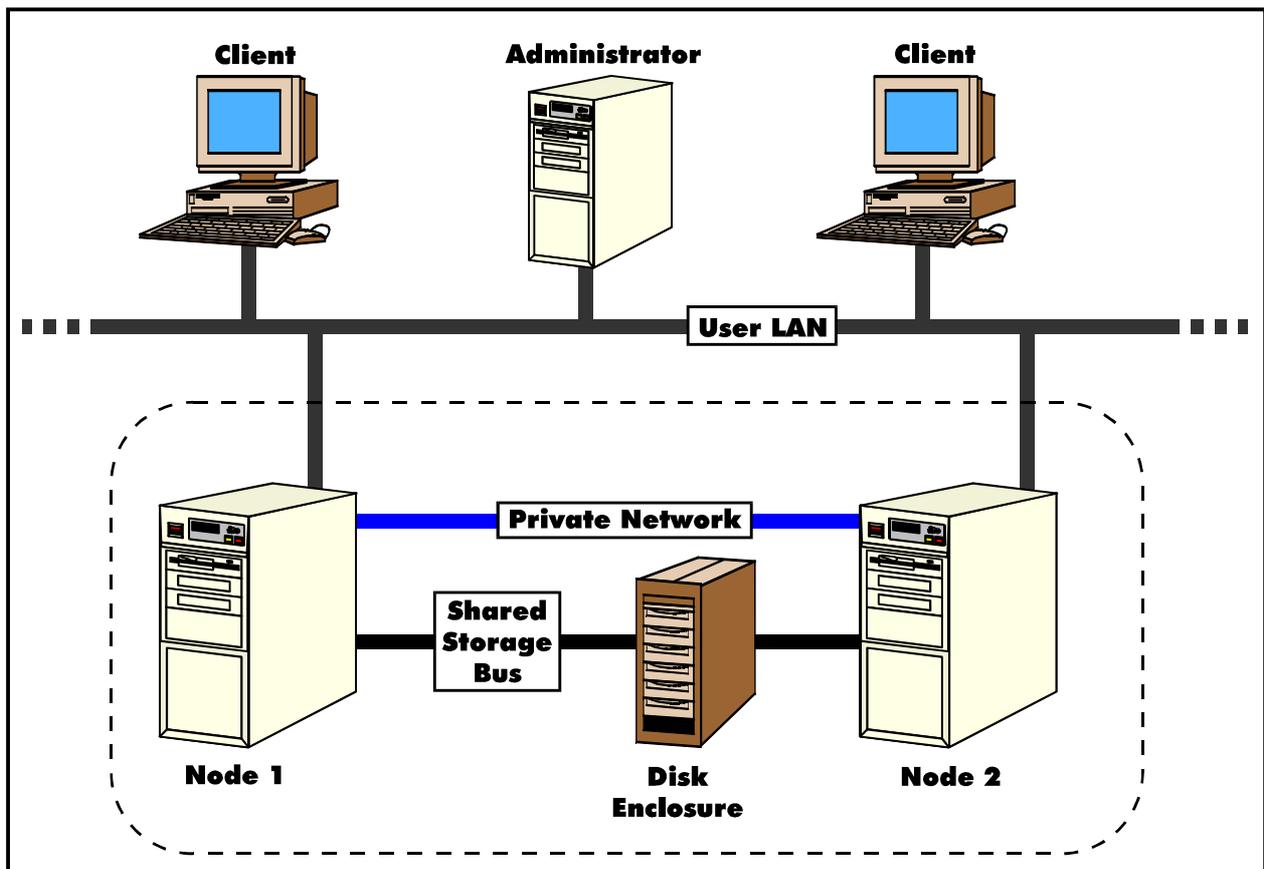


B.5 The Microsoft® Cluster Server Concept

A Cluster Server is always made up of 2 individual NT/2000 Servers, so-called **Nodes**. Both of these Nodes should at least be linked to a commonly used **Shared Storage Bus**. Ultra2 SCSI, Ultra 160 SCSI and Fibre Channel Controllers are suitable for these kinds of links as these technologies enable the configuration of sufficiently long I/O bus systems.

There must be a **Quorum Resource** on the shared storage bus. This can be an own drive, a partition or a sub-directory. This resource is a kind of information-exchange area which will be used by both Nodes and possible running applications on them. An MSCS can not be configured or started without this kind of Quorum Resource. In addition to this resource, other or more disk arrays are normally configured as shared user data hard disks. These disks must be NTFS formatted.

If one of the Nodes fails, the remaining server can take over its tasks. The application which is being carried out on the failed disk is then re-started on the second Node. The applica-



tion has to be installed on both (server) systems. Possible temporary data or information are stored on the Quorum Resource. Corresponding cluster supporting applications can carry on the tasks from the point where the failed Node has stopped via these data and a corresponding API (Application Programming Interface). Other applications can simply only be started anew.

The whole process (take over of tasks for the other Nodes) is classed as a **Failover** by Microsoft.

Furthermore, applications can operate on various Nodes, ensuring that all available resources are used to the optimum. An application can not run on both Nodes at the same time (**Load Balancing**). Microsoft therefore defined this as **Static Load Balancing**, i.e., a static distribution of resources. Dynamic load balancing is not possible at the moment. A speed advantage is only possible if, for example, two different applications are running parallel on different Nodes of a cluster server (e.g., a Web Server on Node 1 and an SQL database server on Node 2).

Both Nodes can communicate with each other via the normal network (User LAN). It is possible to link the Nodes with an additional network (Private Network).

Microsoft urges the user to use two **Interconnects** (1 Public Network as User LAN and a Private Network for the Nodes) for OEMs with original MSCS equipment. The diagram on the previous page shows this kind of configuration.

Data which are on the hard disks of the Shared Storage Bus have to be protected by hardware RAID controllers.

A software RAID solution (e.g., the NT-build-in) is not supported. Microsoft recommends using suitable hardware RAID controllers for protecting data. It is important for these controllers that the RAID configuration data is stored on the disks and not in an NVRAM. This guarantees that even if a controller fails, or after a controller has been exchanged, the RAID configuration still remains intact.

The current ICP Controllers support and are certified (GDT7xy8RN) for operation with Windows NT and Windows 2000 Cluster Server.

B.6 Notes on Novell Netware Orion Clustering

B.6.1 Hardware Requirements and Settings

The following ICP Controller families are able to support **Novell Cluster Services (NCS)**:

- GDT76xxRN
- GDT86xxRZ

It is very important that the ICP Controllers are equipped with a RAIDYNE firmware which supports Clustering. From version **2.30.00** on and upwards support for NCS is included in the Cluster-RAIDYNE firmware. (Cluster-RAIDYNE versions with lower numbers will only support Microsoft Cluster Server).

For the hardware setup the following points should be carefully observed:

1. **All** ICP Controllers which control shared Cluster Drives, have to be connected with each other on **all** IO-channels (SCSI or FC-AL) which host shared drives.

Example 1: Two ICP Controllers control a shared RAID5 Array which consists of three hard disk drives. These three drives are on the same IO-channel. In this example both ICP Controllers have to be connected with this IO-channel.

Example 2: Two ICP controllers control a shared RAID 5 Array which consists of four hard disk drives, where two drives are connected with IO-channel A and the other two with IO-channel B. In this example both ICP Controllers have to be connected with both IO-channels.

2. We highly recommend to use external high quality hot pluggable disk enclosures. These units should support SAF-TE or SES.
3. Shared connections to the enclosure(s) should be secured with SCSI Wonder (in case of shared SCSI channels) or ICP Fibre Hub 1016 (in case of shared Arbitrated Loops).

B.6.2 Install Cluster Array Drives

The boot system and operating system should be isolated from the shared storage. You could think of a separate RAID 1 protected Array Drive.

To setup a Cluster Array Drive which will contain NCS Cluster volumes, follow these rules:

1. With the exception of one server unit, all other server units of the Cluster should be switched off. On this server unit the Cluster Array Drive receives its initial setup.
2. At system boot level of this one server unit press <CTRL>+<G> to enter the rom-resident ICP RAID Console (ICPCON) program.
3. Choose the "Configure Host Drive" menu. At a certain point of the setup procedure ICPCON will ask you if the newly created Host Drive should have the type "Cluster". Confirm with "Y" for yes.
4. If you are not sure if the hard disk drives you use will protect their cache from data loss in case of a Bus reset, switch off the writes caches of these drives. Use the "Configure Physical Devices" menu within ICPCON.

5. Cluster Array Drives of ICP Controllers (from Firmware 2.30.00 on and upwards) are protected against any kind Controller failure with a "Transaction Tracking" system. This system is activated by default. If this feature is not desired (e.g., because of performance reasons), it can be switched off. To accomplish this, all server units of the cluster except one have to be switched off.

B.6.3 Installation of Netware OS and Cluster Services

Please make sure that you install the latest versions of Netware 5.1, the corresponding Service Packs and the Netware Cluster Services. If you are not sure which versions you have and whether they are up-to-date or not, please contact your local Novell distributor or visit Novell's website (www.novell.com).

In any case simply follow step by step the *Installation Notes for the Netware Cluster Services* from Novell. (Skip all steps in this documentation which refer to a Compaq Disk Array – you have already carried out these steps before.)

Important: Even though it will take more time: Never try to install several servers simultaneously. From our experience there is a very high probability that the Cluster will not work in the end. Thus, fully install server by server (except the NCS, which is installed later from a Client). After each installation restart the server. After this restart switch on the next server and begin with the installation there, etc..

Please make sure that the AUTOEXEC.NCF files of all servers of the Cluster contain the following lines:

```
...
LOAD TCPIP.NLM FORWARD=NO
...
NSS /AUTODEACTIVATEVOLUME=ALL
...
```

You should also change the Error Handling settings (e.g., with MONITOR.NLM) into:

```
AUTO RESTART AFTER ABEND = 0
```

B.6.4 The Netware Client

Active-Active Failover capability under Windows NT and Windows 2000 requires an application (e.g., SQL server) which is designed for that scenario. Thus, a standard application (e.g., a simple copy job) will fail and terminate with an error message in an active-active failover situation.

For Windows 95/98 and a Netware Client with version ≥ 3.21 , Active-Active Failover operation with a standard application is possible with some limitations. First, you have to change in *Advanced Settings* of the Netware Client the settings:

Auto Reconnect Level	3
Auto Reconnect Timeout	Default Value (120)
Handle Net Errors	On
Name Cache Level	0
Net Status Timeout	60 (or 120)
Netware Protocol	NDS

(Without these settings massive problems are very likely to occur).

Second, install on the Client *ConsoleOne* – a Java Tool – (current version) to be able to manage the Netware Cluster.

(Important Note: *ConsoleOne* running on the server does not support the Netware Cluster Services).

With *ConsoleOne* you can also install the *Volume Resources*, which appear as Netware Cluster Volumes.

Note: Novell Clustering makes intensive use of the IP protocol. Special care should be taken that all IP addresses (Servers, Clients, Cluster Resources, etc.) are unambiguous.

B.6.5 Information on the Active-Active Failover Scenario

Please be aware that the Netware Cluster Services (as well as the solutions for MS and Linux products) do not represent fully mirrored systems. Thus data loss can happen in case of a server failure with subsequent failover to a backup server.

B.6.6 The ICP Netware Driver and the Tool ICPCON.NLM

Basically, every released version of the ICP HAM driver (GDTRxxxx.HAM) is suitable for Clustering as there are no specific cluster functions in the driver. Though we recommend to use for any new installation the latest released version. Our website contains a download area, where you can download latest versions of all kind of drivers, tools and firmware.

The ICP administration tool ICPCON.NLM allows the administration of the cluster arrays. Please note that a cluster array can only be managed and serviced from one server. Access to this cluster array from other servers (with ICPCON running) is denied. Normally, this is the server in the cluster which was started first and which has found the cluster array before the other servers. If this specific server fails all administration rights are migrated to another, valid server. A server console message notifies the server administrator of this event.

B.6.7 Restrictions

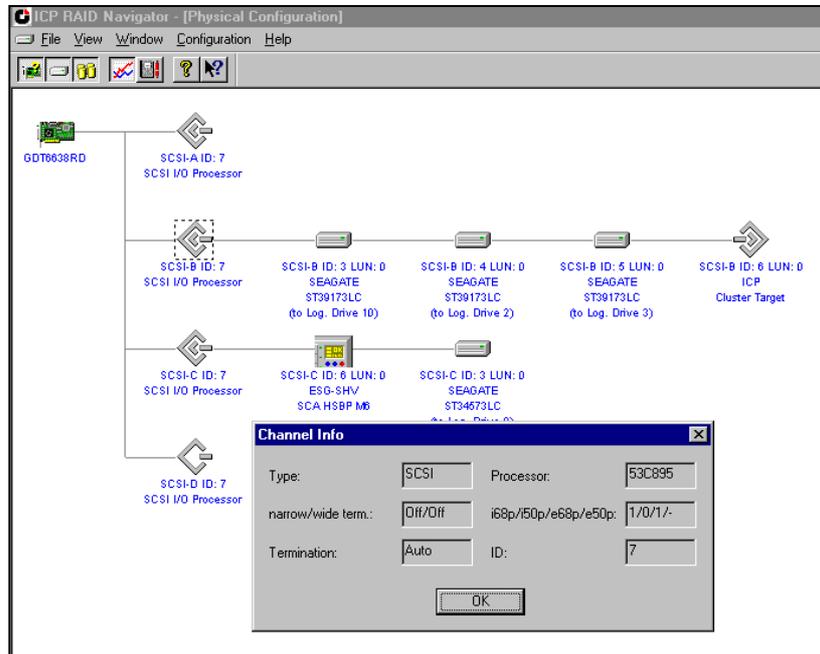
Cluster RAIDYNE in version 2.30 has the following restrictions:

1. Cluster Arrays cannot be expanded online. The Cluster has to be shut down (all servers) and the following steps have to be observed:
 - Shut down the Cluster with all its individual servers and switch them off
 - Switch on only one of the servers and enter ICPCON with <CTRL>+<G>
 - Change the drive type of the Cluster Array into "Private"
 - Expand and/or migrate the array and wait until this process is completed
 - Change the drive type of the Cluster Array back into "Cluster"
 - Start the Cluster again
2. The RAID Levels 1, 10, 15 are currently not supported for Cluster Arrays.

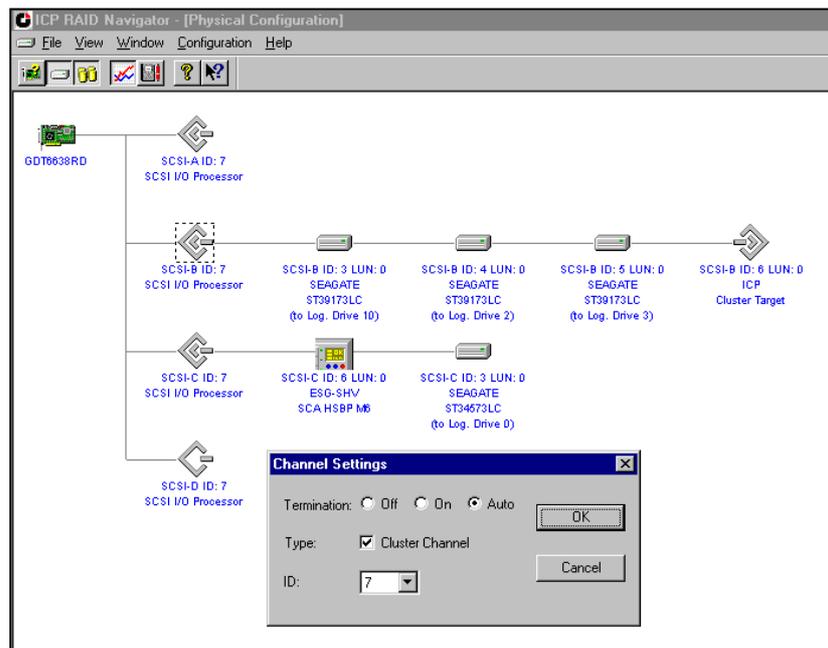
B.7 Notes on the ICP RAID Navigator (ICPRNAV)

A detailed description of the ICP RAID Navigator is given in Chapter D of this manual.

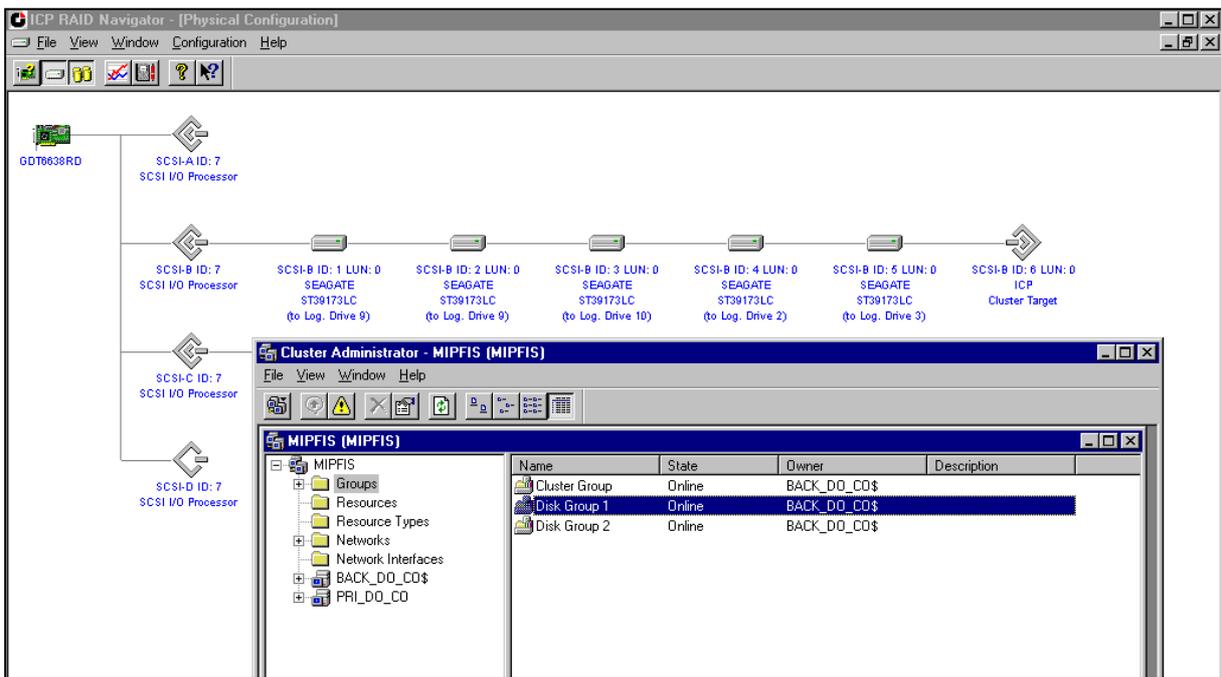
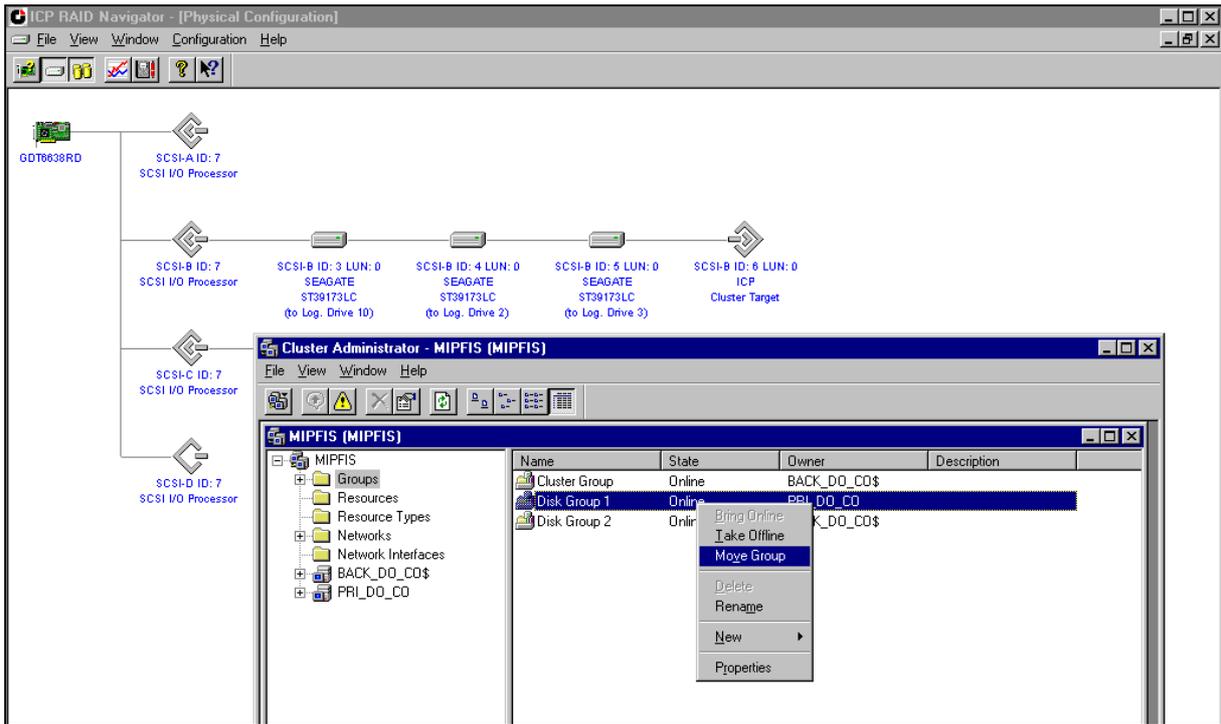
Display of shared I/O channels within ICP RAID Navigator:



This example shows a shared Wide/Ultra2 SCSI channel with five Wide/Ultra2 SCSI hard disks. Two of these hard disks are set up as a RAID 1 disk array and are currently mounted on the GDT6638RD of the other node and are thus not visible here. You could use the Cluster Administrator for Windows NT/2000 to move this resource back to the displayed GDT6638RD. Then, they would also be displayed here.

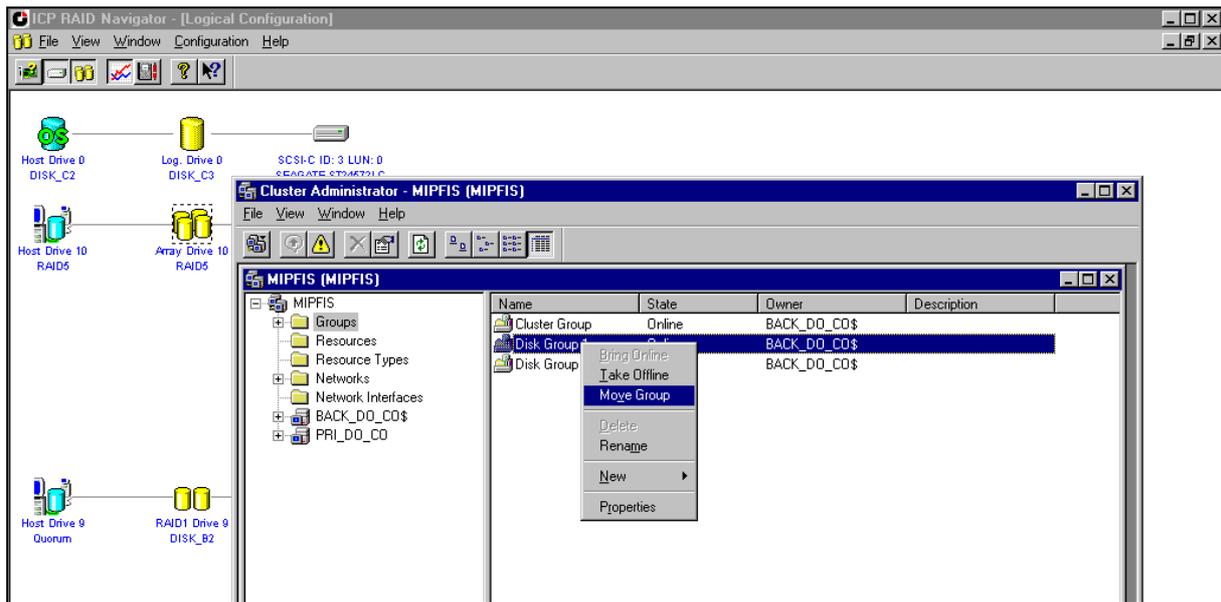


The two disk drives within the RAID 1 Disk Array are moved to the currently displayed ICP Controller:

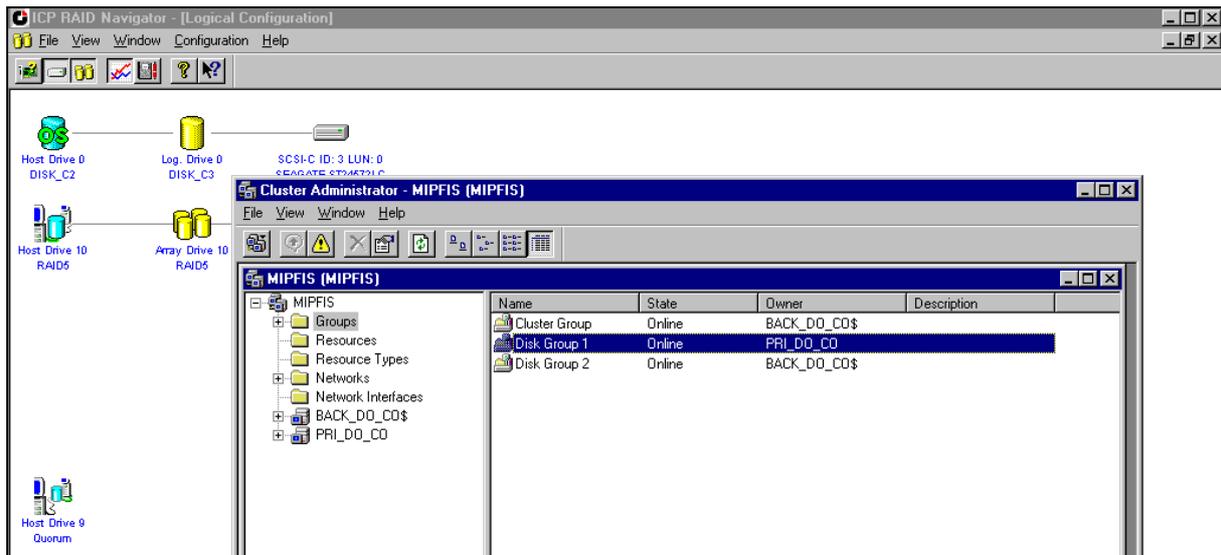


At logical level, the reversed procedure (moving this back to the other ICP Controller) is displayed as follows:

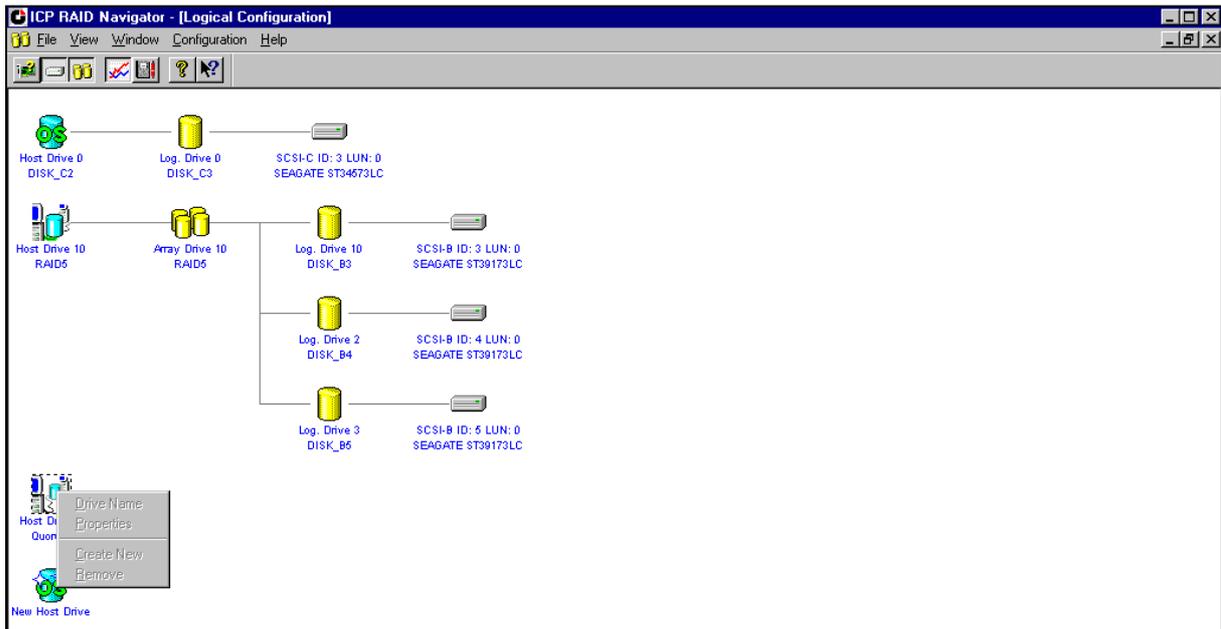
The shared Host Drive still belongs to this ICP Controller. Now it is moved within the Cluster Administrator to the other Node (with the other ICP Controller):



The Host Drive Icon changes showing that this Host Drive belongs to the other Node.

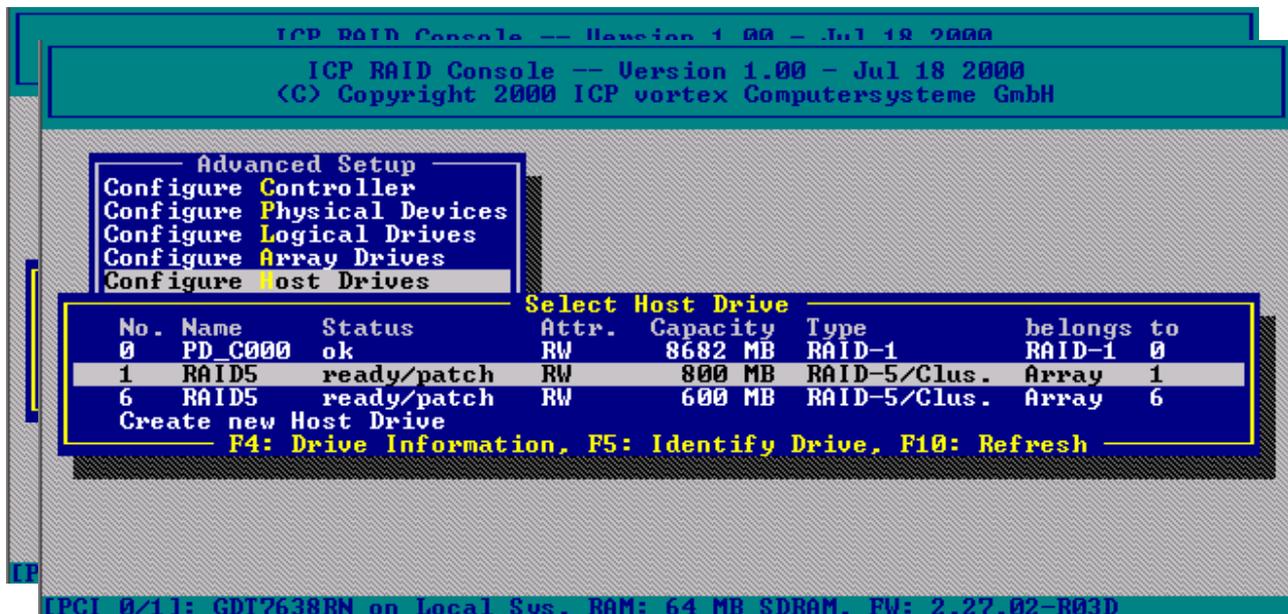


No changes or modifications can be made on this Host Drive. If you loaded the ICP RAID Navigator on the other Node, you would see the exact reversed situation. The RAID 5 Array Drive Icon would stand alone, showing that this Host Drive belongs to the other Node, and the RAID1 Host Drive would be displayed with all its components, allowing all kinds of changes and modifications.



B.8 Notes on ICP RAID Console (ICPCON)

A complete description of ICP RAID Console can be found in the following chapter C. Within the ICP RAID Console all Host Drives, which are shared in a Cluster are marked with the extension „Clus.“.



Chapter C

ICP RAID Console
ICPCON

C. The Program ICPCON

ICP RAID Console (ICPCON) is an extremely helpful and flexible setup and diagnosis tool for the configuration, monitoring, maintenance and tuning of mass storage subsystems which are based on one or more ICP Controllers. Different to the **ICP RAID Navigator** (a GUI-style application for Windows 9x/NT/2000), ICPCON's user interface is character-oriented and available for all operating systems. (Information on the ICP RAID Navigator can be found in a separate chapter of this User's Manual.). In addition to that ICPCON is also part of the ICP Controller's Flash-RAM and can be loaded at system boot level by pressing <CTRL><G>.

Some of the key features of ICPCON:

- ICPCON includes both, diagnosis and configuration functions
- ICPCON is available as both, an executable program under various operating systems, and rom-resident program loadable with <CTRL>-<G> at system boot level (i.e., without any operating system)
- Host Drives can be setup and configured under normal operation
- Online Capacity Expansion of existing Disk Arrays
- Enhanced repair functions for Disk Arrays with failed drives
- Monitoring functions indicate the performance of the various components, Cache statistics
- Available for MSDOS, NetWare 3.x/4.x/5.x, Windows 9x, Windows NT, Windows 2000, Linux, SCO Unix
- Loadable locally (on the server) or remotely from an authorized workstation (support of various protocols); NetBIOS, IPX/SPX, TCP/IP.

C.1 Loading ICPCON

As mentioned before, the ICPCON 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).

C.1.1 Loading the ICPCON Program Under NetWare

The ICPCON program for NetWare is part of the ICP System CDROM. ICPCON can be used either under NetWare 3.x, NetWare 4.x, or NetWare 5.x. There are two different methods of loading ICPCON:

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

Loading ICPCON on the fileserver. Beforehand, the ICP NetWare driver and the auto-loading module CTRLTRAN must have been loaded on the fileserver.

LOAD ICPCON ENTER

on the fileserver.

Loading ICPCON on a workstation. In this case, too, the ICP NetWare driver and the auto-loading module CTRLTRAN must have been previously loaded on the fileserver console. In addition, the module CTRLIPX.NLM has to be loaded. This module searches for a file named CTRLIPX.CFG. This file must be located in the same directory as CTRLIPX.NLM.

The system administrator has to set up a user group named ICP_OPERATOR. All users belonging to this group are given access (through ICPCON) to the ICP Controller(s) in this specific fileserver (Access level 0). Now, the ICPCON program can be loaded from one (or more) workstation(s):

```
ICPCON ENTER
```

C.1.2 Loading the ICPCON Program Under Solaris 7

The ICPCON program for Solaris 7 is part of the ICP System CDROM. To load the program under Solaris 7, enter:

```
ICPCON ENTER
```

C.1.3 Loading the ICPCON Program Under Windows NT / 2000

The ICPCON program for Windows NT / 2000 is part of the ICP System CDROM. To load the program under Windows NT / 2000, enter:

```
ICPCON ENTER
```

For using ICPCON to monitor the server(s) remote, the MON4SOCK.DLL has to be loaded in addition (must be located in the same directory as ICPCON). It supports SPX/IPX and TCP/IP network protocols (for NetBIOS you can load instead MON4NETB.DLL).

C.1.4 Loading the ICPCON Program Under Windows 95/98

The ICPCON program for Windows 95/98 is part of the ICP System CDROM. To load the program under Windows 95/98, enter:

```
ICPCON ENTER
```

For using ICPCON to monitor the server(s) remote, the MON4SOCK.DLL has to be loaded in addition (must be located in the same directory as ICPCON). It supports SPX/IPX and

```
ICPCON ENTER
```

C.1.5 Loading ICPCON Under SCO UNIX

In order to be able to use the ICPCON 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 ICPCON, 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 ICPCON program itself is copied during the SCO UNIX installation into the */etc* directory. ICPCON is loaded by entering:

```
icpcon ENTER
```

C.1.6 Loading ICPCON Under LINUX

The ICP System CDROM includes two archives:

```
icpcon.tgz          ICPCON and object files (intel)
icpcona.tgz        ICPCON and object files (alpha)
```

These archives include all object files to create ICPCON, as well as an executable ICPCON compiled on a current Linux version. If you encounter problems with this executable ICPCON, you can easily compile a new ICPCON on your own Linux system:

```
unpack the tgz-file:      'tar xvfz icpcon.tgz'
compile ICPCON:          'make'
start ICPCON:            './icpcon'
```

In order to be able to compile ICPCON you need the C-compiler and the Kernel sources on your system. The link */usr/src/linux* has to point to the Kernel sources which correspond with the currently booted Kernel of your system. This is important for "signature.c" to use the right magic for the communication with the driver. Otherwise it may happen that you get "Wrong signature" when trying to start ICPCON. ICPCON is loaded by entering:

```
icpcon ENTER
```

C.2 The ICPCON Program

As mentioned before, the ICPCON program appears identical for all operating systems. Thus, 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 ICP firmware. We have defined 4 different levels of hierarchy: Level 1 where the physical devices named **Physical Devices** 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, ICPCON has various menu options, each referring to one level of hierarchy

Host Drives	_____	Level 4
Arrays Drives	_____	Level 3
Logical Drives	_____	Level 2
Physical Devices	_____	Level 1

The following summary gives you an overview of all Host Drive types you can create with the ICP Firmware. **The ICP Controller can simultaneously control several Host Drives of most various types.** For instance, MS-DOS drive C could be a Host Drive of the type disk (consisting of a single hard disk), MS-DOS drive D is a type RAID 5 Array Drive, MS-DOS drive E is a Host Drive of the type chain, and MS-DOS drive F is a CD-ROM which communicates with MS-DOS through corelSCSI and the GDT ASPI manager.

Type of Host Drive	Description of Host Drive	Minimum number of hard disks
Disk	1:1 assignment: Host Drive to hard disk (sometimes also called JBOD)	1
Chain	Concatenation of several hard disks	2
RAID 1	Mirroring	2
RAID 0	Data Striping	2
RAID 4	Data Striping with parity drive	3
RAID 5	Data Striping with striped parity	3
RAID 10	Combined RAID 0 and 1	4

After loading ICPCON, the following screen appears (by pressing F10 you may toggle between black and white display or colored display).

C.2.1 Select Interface

„Interface“ represents the method of how ICPCON connects with the ICP Controller. In this example „Win. 9x/NT/2k“ means that this ICPCON is the Windows version and is loaded on a Windows system.

If you press ENTER, ICPCON scans this local system for ICP Controllers. Instead of „Win. 9x/NT/2k“ you could also find here „Linux“, „NetWare“, etc.

Selecting „Sockets“ allows for the remote connection of this workstation with servers using ICP Controllers. You can either choose TCP/IP or IPX/SPX network protocols. In order to access the servers with ICP Controllers, the corresponding remote service must be loaded on the server. In addition the supervisor has to setup users and assign access rights (Name, password, see also chapter K.9)

The screenshot shows the ICP RAID Console interface. At the top, a blue header displays "ICP RAID Console -- Version 1.00 - May 30 2000" and "(C) Copyright 2000 ICP vortex Computersysteme GmbH". Below this, a central menu titled "Select interface" is displayed on a dark background with light text. The menu options are "(win. 9x/NT/2k)", "Sockets", and "IBM NetBIOS". A yellow prompt "F10: Toggle colors" is visible at the bottom of the menu. At the bottom of the screen, there are two blue input fields. The left one is titled "Remote Machine" and contains "Name/Address: 192.89.3.35". The right one is titled "192.89.3.35 login" and contains "Password: *****".

After selecting „Sockets“ and „TCP/IP“ you may enter the IP address of the server (if you would have chosen „SPX/IPX“ ICPCON would scan the network for suitable servers, which have SPX/IPX protocol). After that you may enter your user name and password.

C.2.2 Select Controller

After this login procedure ICPCON delivers a list of ICP RAID Controllers which are installed in this server (in this example one GDT7563RN). This list contains information on the controller name, the PCI Bus system (0=primary, 1=secondary, etc.) and separated with a slash the PCI slot number, the controller's features (C=Chaining, 0=RAID 0, 1=RAID 1, 4=RAID 4, 5=RAID 5, 10=RAID 10) and the firmware level. After selecting the controller with ENTER, all further settings and changes to these settings within ICPCON refer to this ICP Controller and the connected devices.

Select Controller				
No.	Name	Position	Features	Firmware
0	GDT7563RN	[PCI 0/1]	C 0 1 4 5 10	2.27.00-RFFF

F10: Toggle Colors

5=RAID 5, 10=RAID 10) and the firmware level. After selecting the controller with ENTER, all further settings and changes to these settings within ICPCON refer to this ICP Controller and the connected devices.

C.2.3 The two Menu Areas „Monitor“ and „Express/Advanced Setup“

ICPCON offers two fundamentally different operating modes:

- Express Setup / Advanced Setup with configuration functions
- Monitor with monitoring functions

Express Setup	
Configure Host Drives	
Repair Array Drives	
Monitor	
View Statistics	
View Events	
View Hard Disk Info	
Save Information	
F4: Advanced Setup	

Advanced Setup	
Configure Controller	
Configure Physical Devices	
Configure Logical Drives	
Configure Array Drives	
Configure Host Drives	
Repair Array Drives	
Monitor	
View Statistics	
View Events	
View Hard Disk Info	
Save Information	
F4: Express Setup	

The various menu options can be selected either with the cursor up/down keys, or by pressing the high-lighted character. F4 allows the switching between the Advanced Setup and Express Setup modes.

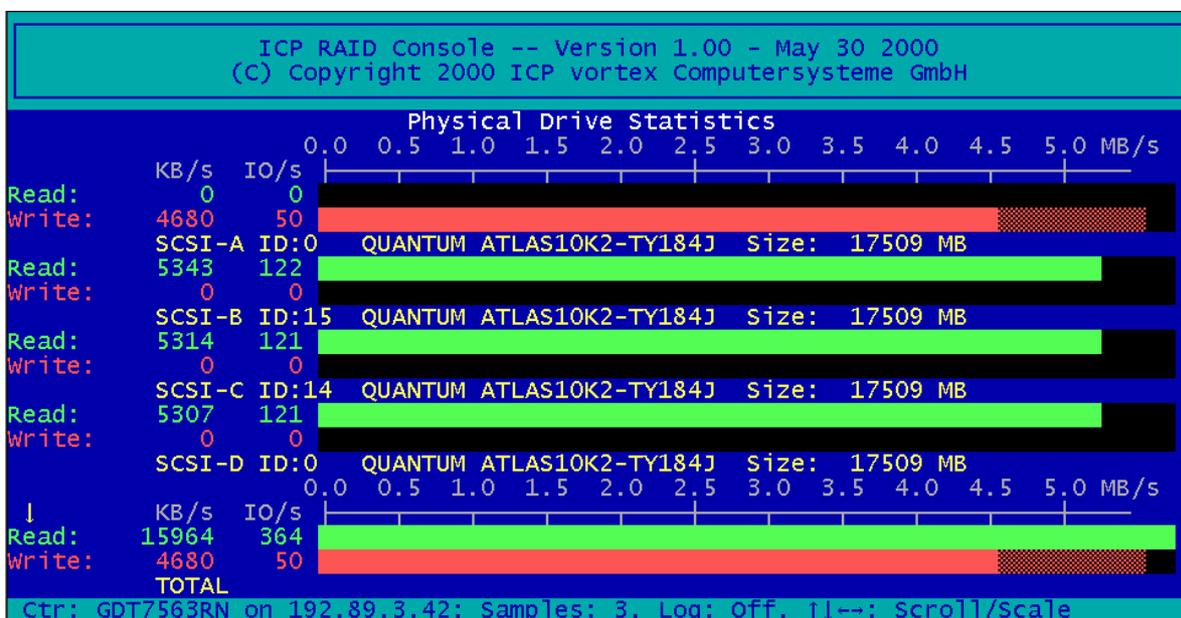
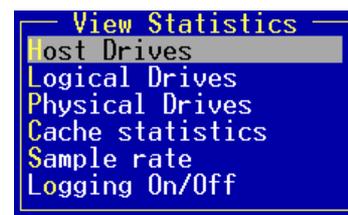
Monitor	Express Setup / Advanced Setup
<ul style="list-style-type: none"> ▪ With „View Statistics“ you may supervise the performance of the different components of the RAID subsystem. ▪ „View Events“ is an enhanced event recorder for RAID and controller specific events. ▪ With „View Hard Disk Info“ you can view detailed information on the connected devices (important are the retries, reassigns, grown defects and last status information) 	<ul style="list-style-type: none"> ▪ „Configure Controller“ allows you to setup the ICP Controller. Here, you can change the Cache settings, the termination, the memory test, etc. ▪ With „Configure Physical Devices“ you can initialize the devices at physical level, e.g., change the SCSI protocol, transfer rates, etc. ▪ „Configure Logical Drive“ allows you to configure or create Logical Drives.

- „Save Information“ creates a complete protocol file of the current RAID sub-system including all settings of the ICP Controller and the drives. This file can be used for documentary reasons or for remote diagnosis.
- „Configure Array Drives“ allows you to configure or create Array Drives.
- With „Configure Host Drives“ you can either configure already existing Host Drives, or create new ones. I.e., you can create under normal operation new Host Drives without shutting down the computer.
- The „Repair Array Drives“ menu offers very powerful menu-guided functions to repair Array Drives which have failed drives.

C.3 The Menu Monitor

C.3.1 Menu Monitor: View Statistics

In each of these menus, ICPCON displays the performance of the drives of the corresponding level of hierarchy. The throughput of the drives is displayed in KB/s (KiloBytes per second) and IO/s (I/Os per second, number of IO's on the controller). The performance figures reflect the load being on the controller and not necessarily the maximum performance the controller can deliver.



After selecting one of the Drives Statistics ICPCON displays a list of all Drives of this level (for this example the Physical Drives, i.e., level 1). In addition to the performance report on the Drives, you are given additional information on each device. The ICP I/O channel the hard disk is connected to, which ID the hard disk has, the name of the hard disk, the gross

capacity (1MB = 1024KB). 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 Drives (if available).

After selecting the menu option „Cache Statistics“ you can view the utilization of the ICP Controller's Caches, separated in the Read Cache and the Write Cache. This menu also displays the size of the Cache in KB and the settings of both Caches (On, or Off). The figures for „Cache Hits“ show the how often requests can be serviced out of the cache, i.e., without triggering an immediate Disk IO.

By setting the „Sample 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.

Logfile: `hoststat.log`

With „Logging On/Off“, you may create a log file which records all the statistic values over a longer period. If you choose On, ICPCON asks for the path/name of the log file.

C.3.2 Menu Monitor: View Events

With „View Events“ ICPCON displays all ICP Controller Events. They can also be recorded and saved into a log file. This function gives the administrator a good help to analyze and supervise ICP Controllers with Array Drives.

Controller Events						
Ctr	Time	Cnt	Evt	Text		
0	Thu Jan 01 01:00:00 1970	1	20	Array Drive 3: parity build finishe...		
0	Thu Jan 01 01:00:00 1970	1	42	Array Drive 3: parity build started		
0	Thu Jan 01 01:00:00 1970	1	56	Host Drive 7 created		
0	Thu Jan 01 01:00:00 1970	1	52	Array Drive 3: expand finished succ...		
0	Thu Jan 01 01:00:00 1970	1	51	Array Drive 3: expand started		
0	Thu Jan 01 01:00:00 1970	1	20	Array Drive 3: parity build finishe...		
0	Thu Jan 01 01:00:00 1970	1	42	Array Drive 3: parity build started		
0	Thu Jan 01 01:00:00 1970	1	60	Array Drive 3: parity build aborted		
0	Thu Jan 01 01:00:00 1970	1	42	Array Drive 3: parity build started		
0	Thu Jan 01 01:00:00 1970	1	20	Array Drive 3: parity build finishe...		
0	Thu Jan 01 01:00:00 1970	1	42	Array Drive 3: parity build started		

↑↓: scroll, F8: clear, F2: save + clear

C.3.3 Menu Monitor: View Hard Disk Info

This menu shows for all Drives :

ICP RAID Console -- Version 1.00 - May 30 2000 (C) Copyright 2000 ICP vortex Computersysteme GmbH							
Hard Disk Information							
Chn	ID	Vendor	Product	Retr./Reass.	Grown	Def.	Last Status
SCSI-A	0	QUANTUM	ATLAS10K2-TY184J	0/0	0	0	0x00000000
SCSI-B	15	QUANTUM	ATLAS10K2-TY184J	0/0	0	0	0x00000000
SCSI-C	14	QUANTUM	ATLAS10K2-TY184J	0/0	0	0	0x00000000
SCSI-D	0	QUANTUM	ATLAS10K2-TY184J	0/0	0	0	0x00000000
SCSI-E	1	IBM	DPSS-309170M	0/0	0	0	0x00000000
SCSI-E	3	IBM	DPSS-309170M	0/0	0	0	0x00000000
SCSI-E	6	IBM	DDYS-T18350M	0/0	0	0	0x00000000
SCSI-F	6	QUANTUM	ATLAS10K2-TY184J	0/0	0	0	0x00000000

- the SCSI channel
- the SCSI-ID
- the vendor and type
- Retries/Reassigns, Grown Defects and the Last Status

(1) The *Retries* counter is incremented by one unit whenever the ICP Controller retries to access a hard disk. If this counter continues to increase (possibly on other hard disks, too) it is very likely that the 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 hard disk concerned defective. The retry counter also increases when the SCSI parameters of a hard disk are changed (see further ahead). Obviously, retries due to this do not imply bad cabling.

(2) The reassign counter reflects the number of media defects which occur on the hard disk drive. Defective blocks of the hard disk 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 hard disk through according reassignment tables. Note: If a hard disk 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.

(3) The Grown Defects counter shows the number of media defects which have occurred since the first time the device was operated with an ICP Controller. A specific 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.

(4) The Last Status information should always be 0x00000000. After a device failure or other significant events, a different value may be displayed here. This value is volatile and is reset to 0x00000000 after each power up and/or reset.

If you press ENTER on a Drive, ICPCON display further information on the SCSI parameter settings.

```

Hard disk Information
Name .....: QUANTUM ATLAS10K2-TY184J
Capacity .....: 17509 MB

Sync. Transfer .....: Enabled
Sync. Transfer Rate : 160.0 MB/s
Sync. Transfer Mode.: DT
Data Path Width ....: Wide
Domain Validation ..: On
Disconnect .....: Enabled
Disk Read Cache ....: On
Disk Write Cache ...: On
Tagged queues .....: On

Grown defects .....: 0
Last status .....: 0x00000000

```

C.3.4 Menu Monitor: Save Information

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

At the end of this protocol is a chronological listing of boot messages and other events stored in the Flash-RAM of the ICP Controller. If the buffer is full, the oldest events are deleted first.

```

View Information
ICP RAID Console -- Version 1.00 - May 30 2000
=====
Date/Time:      Tue Jun 13 11:27:41 2000
Operating system: Windows 9x/NT/2000, Version: 5.00, Rev. 2195
Processor:      x86 Family 6 Model 7 Stepping 2
Physical memory: 130596 KB

Controller Information
=====
Type:           GDT7563RN      Slot/Address:    0/1
Firmware version: 2.27.00-RFFF      Serial no.:      FFEFFFFFFF
Features:       C 0 1 4 5 10  Channel count:    6
Cache:         On           Delayed write:   On
Hardware level: 0           Cache RAM:      64 MB, 2 bank(s)

```

In the menu „Configure Controller“ the logging buffer can be cleared with „Clear Log Buffer“. The logging list is a good source to analyze complex events and problems.

```

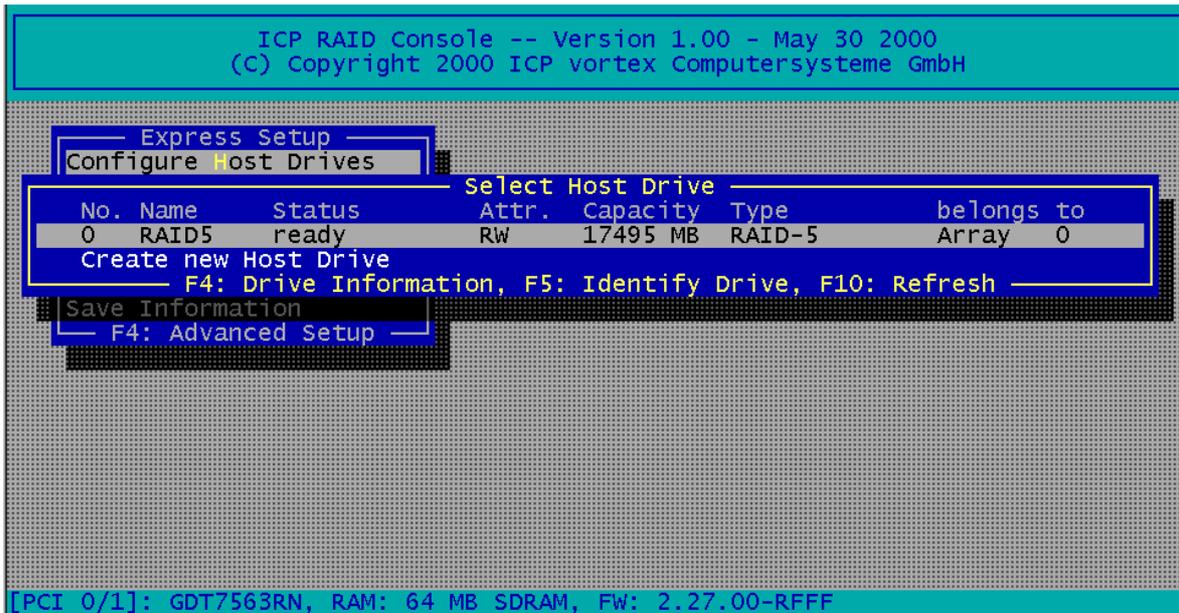
View Information
Screen service messages / Async. events
=====
Tue Jun 13 14:16:52 2000: Array Drive 3: parity build finished successfully
Tue Jun 13 14:13:34 2000: Array Drive 3: parity build started
Fri Jun 09 13:36:25 2000: Host Drive 7 created
Fri Jun 09 13:36:25 2000: Array Drive 3: expand finished successfully
Fri Jun 09 13:24:59 2000: Array Drive 3: expand started
Fri Jun 09 13:24:36 2000: Array Drive 3: parity build finished successfully
Fri Jun 09 13:21:42 2000: Array Drive 3: parity build started
Fri Jun 09 13:21:18 2000: Array Drive 3: parity build aborted
Fri Jun 09 13:19:34 2000: Array Drive 3: parity build started
Fri Jun 09 13:18:25 2000: Array Drive 3: parity build finished successfully
Fri Jun 09 13:14:45 2000: Array Drive 3: parity build started
Fri Jun 09 11:39:24 2000: RAID-5 Host Drive 3 installed (ready)
Fri Jun 09 11:39:24 2000: RAID-5 Host Drive 0 installed (ready)
Fri Jun 09 11:39:23 2000: SCSI-E ID:6 LUN:0 -- IBM      DDYS-T18350M
Fri Jun 09 11:39:23 2000: SCSI-E ID:3 LUN:0 -- IBM      DPSS-309170M

```

C.4 The Menu Express/Advanced Setup

C.4.1 Menu Express Setup: Configure Host Drives

This function allows a very easy installation of new Host Drives and does not require any special knowledge. Apart from minor differences, this menu option is identical with the menu option „Configure Host Drives“ in the Advanced Setup (in Express Setup, the user may not select a stripe size (defaults to 128KB) or use the Split/Merge functions.)



After selecting Configure Host Drives, ICPCON displays a list of already existing Host Drives. With a new system this list will show no entries. In this example there is already one Host Drive available. It's name is "RAID 5" and it is an RAID 5 Array Drive (with approx. 17GB capacity). The status is "ready". I.e., this Host drive is fully available and redundant. The following states for RAID Host Drives are possible:

"Idle" State

This state is characterized by the fact that the redundant information of the disk array has never been entirely created. The disk array is in this state after its first configuration and until you quit ICPCON. If an error should occur while the array is in the *build* state, the array returns to the *idle* state (exception: if during *build* mode the dedicated drive of RAID 4 fails, the mode changes to *fail*).

"Build" State

After the disk array has been configured for the first time, it assumes the *build* state as soon as you quit ICPCON. While the array is in the *build* state, redundancy information is calculated and stored to the hard disks of the array.

"Ready" State

The disk array is fully operational when in the *ready* state. All redundant information is present, that is, a hard disk can fail without impairing the functionality of the disk array. This is the normal state of a disk array. The state *ready/expand* indicates, that the RAID level and/or capacity are currently migrated/expanded.

"Fail" State

The disk array changes to the *fail* state whenever a Logical Drive fails. Redundancy information is still present, thus allowing the remaining hard disks to continue working. This state should be eliminated as soon as possible by replacing the defective hard disk. If a so-called Hot Fix drive has previously been assigned to a disk array with ICPCON, the controller will automatically replace the defective drive and start the reconstruction of the data and the redundant information. Therefore, under these circumstances the *fail* state is only temporary and will be eliminated by the controller itself.

"Rebuild" State

The disk array will assume this state after the automatic activation of a Hot Fix drive or after a manual replacement carried out with ICPCON. The data and the redundant information are reconstructed and stored to the new drive.

"Expand" State

If the capacity or RAID level of an existing disk array is changed, the disk array changes its state into *expand*. As soon as the expansion or migration is completed, the state changes back to *ready*.

"Error" State

If a second hard disk should fail while the disk array is in the *fail* or *rebuild* state, it is not possible to continue the working session without restrictions. The disk array is still available for I/Os, but data loss and error messages on the host level are possible.

The following state diagram of the disk array summarizes the states described above and the transitions from one state to another.

Some of these states may become the addendum **patch** (e.g. *build/patch*, *ready/patch*).

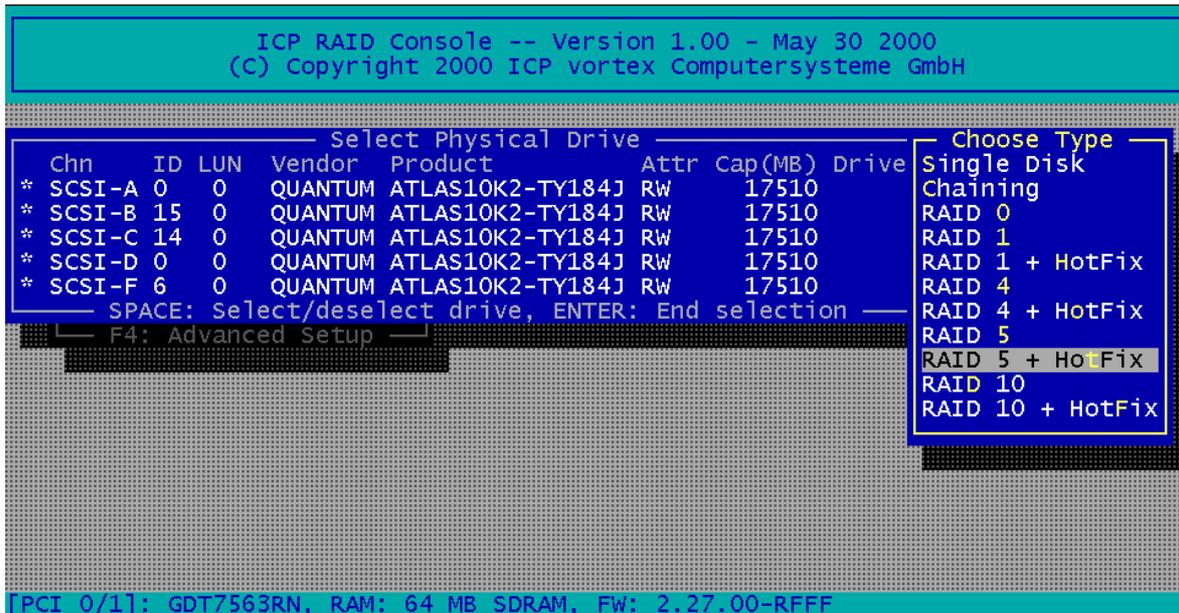
This word indicates that the original Array Drive went through a significant procedure. I.e., the parity information was recalculated anew.

Or, the Array Drive has been patched from the error state into the fail state. This may become extremely helpful in a situation where two Logical Drives of an Array Drive, fail at the same time, but only one of the two Logical Drives is really defective and the other was blocked out, since it was connected with the same SCSI channel as the defective one. The Array Drive's state is error and normally all data would be lost. The ICP Controllers include some functions, which allow the patch of this Array Drive from the error state into the fail state. Before the actual patch, the defective drive has to be physically removed from the Array Drive. Such a patch-procedure is a real sheet-anchor and should only be used, after a detailed consultation with a trained support person (a printout of the *Save Information* file, is extremely helpful).

Pressing F4 delivers level by level detailed information on a Host Drive and its components. This may assist to get an easy overview of a specific Host Drive. After pressing F5 the ICP Controller switches the LEDs (if available) of the hard disks belonging to that Host Drive on and off.

After selecting „Configure new Host Drive“, the ICP Controller scans all channels for free hard disks (i.e., hard disks which are not yet part of a Host Drive) and displays these units in a list. Use the Space bar to select/deselect hard disks and the cursor up down keys to move the selection bar from one hard disk to another. Marked hard disks show an "*" in the first column. In the choose Type windows all possible Host Drive types are displayed (the possible types depend on the number of selected hard disks and the firmware level of the ICP Controller).

Finish the selection by pressing ENTER.



ICPCON displays a security message pointing out that all existing data on the selected hard disks will be destroyed after confirming with Yes.

Do you want to create a host Drive from the selected disk(s) ?
(CAUTION: All data will be destroyed !) (Y/N)

After pressing "Y" the user may limit the capacity per hard disk which will be used for the Host Drive. This can be very helpful for the procurement of future spare hard disks. After

Used Capacity per Drive (1..17509 MB): 17000_

that ICPCON automatically creates and configures the new Host Drive and adds it to the list. All SCSI parameters are adjusted to optimum values.

Select Host Drive						
No.	Name	Status	Attr.	Capacity	Type	belongs to
0	RAID5	ready	RW	17495 MB	RAID-5	Array 0
3	RAID5	build	RW	52528 MB	RAID-5	Array 3

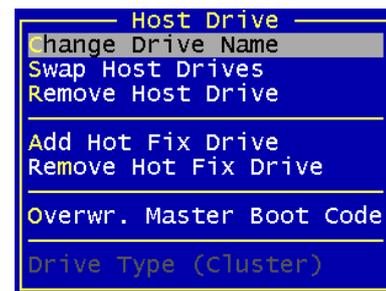
Create new Host Drive
F4: Drive Information, F5: Identify Drive, F10: Refresh

After leaving ICPCON or after pressing F10 (for refresh) the ICP Controller starts the build process on this Host Drive. In our example it calculates the parity information and writes it to the hard disks. The status during this process is build. When the build is successfully finished the status changes into ready.

If you press ENTER on a Host Drive the following options become available:

"Change Drive Name" allows you to change the name of a Host Drive. The name serves to identify a Host Drive with ICPCON. E.g., you could name the boot Host Drive "Boot" and the data Host Drive "Data".

"Swap Host Drives". When the PCI computer is switched on, the Host Drives are initialized in the order of the Host



Drive list, which means that the operating system is booted from the Host Drive having the lowest number. For reasons of flexibility, a Host Drive's position in the list can be changed. However, the position of the Host Drive from which the operating system is booted and the position of the Host Drive from which ICPCON (disk version) was started (both can be the same), cannot be changed. If you wish to change the position of these drives, you have to boot the operating system and ICPCON from a floppy disk or use the ICPCON from the Flash-RAM of the ICP Controller. To change the position of a Host Drive in the Host Drive list, highlight the Host Drive and confirm with ENTER. Then, type on the new position and press ENTER.

“Remove Host Drives”. Removing a Host Drive is a serious action. **All data will be lost after removal.** If you want to remove a Host Drive belonging to an Array Drive for which several Host Drives exist (after capacity expansion, or after splitting), all other Host Drives will also be removed.

“Split Host Drive”. For some purposes it might of interest to split an existing Host Drive into two or several Host Drives. Each Host Drives looks to the operating system just like a single hard disk. Since the new Host Drives have smaller capacities ICPCON has to write new header information on the two Host Drives. **All data will be lost.**

“Merge Host Drives”. This function reverses the *Split Host Drive* option. Only such Host Drives can be merged which belong to the same Array Drive or Logical Drive. Since the new Host Drives has a larger capacity ICPCON has to write a new header information on the new Host Drives. **All data will be lost.**

“Partition Host Drive”. This option is not available, when loading ICPCON from the Flash-RAM of the ICP Controller. Before you can partition a new Host Drive it may become necessary to reboot the system, first. The partitioning menu has similar functions as the MS-DOS program FDISK. You can create and delete a partition and also change the active partition. MS-DOS can only be booted from an active partition. Just like FDISK, ICPCON can handle primary partitions, extended partitions, and logical drives within the extended partitions.

“Add Hot Fix Drive” allows you to add a Hot Fix drive to an existing RAID 1, RAID 4, RAID 5, or RAID 10 Array Drive. There are two different types of Hot Fix drives: *Private* and *Pool* Hot Fix drives. 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 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.

A *Private* Hot Fix drive is dedicated to one RAID 1, RAID 4, RAID 5 or RAID 10 Array Drive.

“Remove Hot Fix Drive” allows you remove a previously assigned Hot Fix drive.

“Overwrite Master Boot Code”. This option creates a valid and consistent master boot record on the selected Host Drive and should be carried out on any new Host Drive on which Windows NT is installed. **Never use this function when the Host Drive contains valid data, all data will be lost.**

The option “Drive Type (Cluster)” is available only with ICP Controllers which are equipped with Cluster RAIDYNE® (GDTx6xxx) and allows to assign one Host Drive to several ICP Controllers (Type Cluster).

C.4.2 Menu Express Setup: Repair Array Drives

This function allows the online repair of Array Drives which show failed drives. After selecting this menu option, ICPCON displays a summary of all installed Array Drives (in this example 2) and the number of Array Drives which are in critical states.

In this example there is one Array Drive in the FAIL state. I.e., the Array Drive is still operating but longer redundant.

After pressing any key, ICPCON displays a list of Array Drives which are candidates for this online automatic repair. Note: Array Drives which have the „ERROR“ state are very critical and have lost 2 or even more drives. These Array Drives cannot not be repaired with this function. In such critical cases the data integrity can longer be maintained. You may call our technical support center for further assistance.

```

Summary
2 Array Drives installed
1 Array Drives in FAIL state
0 Array Drives in ERROR state
0 Array Drives in BUILD state
0 Array Drives in REBUILD state
0 missing Private HotFix Drives

Press any key.

```

After selecting the Array Drive, ICPCON display the actual drive which has failed. In this example it is the hard disk which forms Logical Drive 5.

```

Select Drive to repair
No. Name Type State
2 RAID5 RAID-5 fail

```

After confirmation, the failed drive has to be taken out and the new one has to be configured on the same ID as plugged in again.

```

Monitor
Repair Array Drives
Vi Select Drive to repair
Vi No. Name Type State
Vi 2 RAID5
Sa F4: Config

Physical Disk SEAGATE ST39175LC
at channel SCSI-B, ID 10,
which belongs to Logical Drive 5
failed. Do you want to replace it? (Y/N)

```

```

The Disk SEAGATE ST39175LC at channel SCSI-B, ID 10 will be replaced.
Please plug in a disk at the same channel with the same ID.
YOU MUST NOT DO ANY UNPLUGGING OR PLUGGING ON OTHER CHANNELS

Press any key after hot plug.

```

ICPCON detects the new drive and expects a clearance to build this drive as a replacement into the Array Drive.

```

You plugged in the disk SEAGATE ST39175LC , ID 10 at channel SCSI-B.
Do you want to add this disk to the Array? (Y/N)

```

After that the state of the Array Drive changes into „REBUILD“, i.e. the missing data is reconstructed out of the remaining data and the redundancy information.

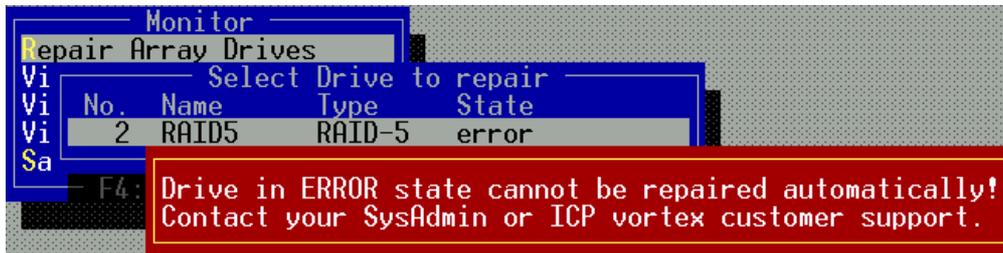
```

Summary
2 Array Drives installed
0 Array Drives in FAIL state
0 Array Drives in ERROR state
0 Array Drives in BUILD state
1 Array Drives in REBUILD state
0 missing Private HotFix Drives

Press any key.

```

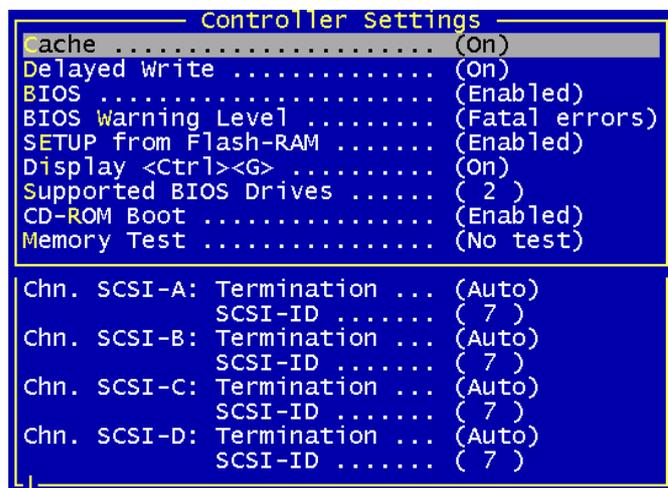
The „ERROR“ state of an Array Drive is very critical. There are several procedures in the ICP Controller's firmware to handle such cases and bring back the Array Drive in operation without losing data. The most suitable procedure for the specific case, should be elaborated with your system administrator or our technical support. (Note: ICP vortex offers 2-day RAID Workshops in our training center, which also focus on the resolution of such problems.



C.4.3 Menu Advanced Setup: Configure Controller

C.4.3.1 Menu Advanced Setup: Configure Controller, Controller Settings

Within the Controller Settings all relevant parameters and settings concerning the ICP Controller itself can be changed. All information concerning Physical, Logical, Array and Host Drives are stored twice (*primary and secondary configuration data*) on each hard disk. I.e., the ICP Controller itself carries no configuration data of a specific disk array. Out of this reason, the migration of Host Drives between ICP Controllers is very easy and secure, even if the SCSI IDs and channels are changed. Simply connect the hard disks to the new ICP Controller, and the Host Drives will automatically be available again. The various parameters and settings have the following meaning:



Parameter	Alternatives	Description
Cache	On	ICP Controller cache (write and read cache) On
	Off	ICP Controller cache (write and read cache) Off
Delayed Write	On	Write cache On
	Off	Write cache Off
BIOS	Enabled	BIOS enabled
	Disabled	BIOS disabled, but <CTRL><G> allowed
	Removed	BIOS and <CTRL><G> disabled
BIOS Warning Level	All Messages	All ICP BIOS warnings and errors are displayed
	Fatal Errors	Only fatal messages are displayed
SETUP from Flash-RAM	Disabled	<CTRL><G> not possible
	Enabled	<CTRL><G> allowed

Display <CTRL><G>	Off On	<CTRL><G> not displayed at system boot level <CTRL><G> displayed at system boot level
Supported BIOS Drives	2 7	The ICP BIOS supports 2 drives under MS-DOS The ICP BIOS supports 7 drives under MS-DOS
CD-ROM Boot	Enabled Disabled	Boot from CD-ROM enabled Boot from CD-ROM disabled
Memory Test	No test Standard Double scan Intensive	No test at all Standard test Double scan test Intensive test (takes longer)
Chn. SCSI-A Termination	Off On Auto	SCSI termination of the channel Off SCSI termination of the channel On SCSI termination according to occupied connectors
SCSI-ID	0,1,2,3,4,5,6,7	Possible SCSI IDs of the channel

C.4.3.2 Menu Advanced Setup: Configure Controller, Firmware Update

The firmware, the BIOS and the ICPCON program of the ICP Controller are stored in a Flash-RAM which is part of the ICP Controller hardware. In contrast to EPROMs, Flash-RAMs can be re-programmed many times and without the complicated UV-light erasing procedure. Firmware, BIOS and ICPCON are part of the **GDT_RPFW** file. The file has an extension (e.g. GDT_RPFW.009) which indicates the version stepping. The latest version of the this file can be downloaded from our Website (www.icp-vortex.com). This menu option is not available when ICPCON is accessing the ICP Controller remotely.

C.4.3.3 Menu Advanced Setup: Configure Controller, Intelligent Fault Bus

Intelligent Fault Bus is an older subsystem standard, which is no longer used in modern subsystems or backplanes. ICP's Ultra/Wide Disk Array Controllers were the last ICP Controllers supporting this, sometimes also called "DEC™ Fault Bus". Today's modern subsystems are either using SAF-TE (SCSI Accessed Fault Tolerant Enclosures) or SES (SCSI Enclosure Services) as communication links to the controller. These intelligent subsystems are normally build on so-called backplanes, which host hard disks which are equipped with SCA connectors (Single Connector Attachments). The backplane has a dedicated electronics with microprocessor and firmware which allows an intelligent communication between the subsystem/backplane and the ICP Controller. Through this channel the subsystem can for example report its temperature and the power supply status to the ICP Controller. The major objective of SAF-TE or SES is to provide the Auto Hot Plug. In contrast to the Hot Plug, a defective drive is simply pulled out the subsystem and the replacement unit is plugged in again. Both, the ICP Controller and the subsystem control this process so that problems on the SCSI/FC-AL bus are eliminated and the rebuild of the missing data is initiated fully automatically. No further user interaction is necessary.

A unique and special feature of ICP Controllers is the following:

C.4.3.4 Menu Advanced Setup: Conf. Controller, Non-Intelligent Enclosures

Since SAF-TE or SES subsystems are pretty cost intensive, ICP has integrated a special functionality which can control up to 16 non-intelligent enclosures with up to 15 hard disk, each, and provides the Auto-Hot-Plug. It is clearly up to the user's responsibility to use components (e.g., disk shuttles), which are really hot pluggable.

Select Enclosure			
No.	Vendor	Name	Status
0	MALIBU	PA46-350	Configured
1			Not Configured
2			Not Configured
3			Not Configured
4			Not Configured
5			Not Configured
6			Not Configured
7			Not Configured
8			Not Configured
9			Not Configured
10			Not Configured
11			Not Configured
12			Not Configured
13			Not Configured
14			Not Configured
15			Not Configured

ENTER: Select, F4: Edit Vendor/Name

In this example one subsystem is already defined. Naturally, the hard disks with their disk shuttles could be also directly mounted in the server enclosure. The term "Enclosure" in these cases is more a definition set, which includes all hard disks which should be auto hot pluggable.

The following example shows the slots of the enclosure which have been assigned with hard disks. To fill an empty slot, press ENTER and select the desired hard disk.

Enclosure Slots								
Slot	Chn	ID	LUN	Vendor	Product	Attr.	Cap(MB)	Drive
0	SCSI-A 0	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 4
1	SCSI-B 15	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 5
2	SCSI-C 14	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 6
3	SCSI-D 0	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 7
4	SCSI-E 1	0	0	i	IBM DPSS-309170M	RW	8747	Drive 0
5	SCSI-E 3	0	0	i	IBM DPSS-309170M	RW	8747	Drive 1
6	SCSI-E 6	0	0	i	IBM DDYS-T18350M	RW	17501	Drive 2
7	SCSI-F 6	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 3
8	No hard disk installed in this slot							
9	No hard disk installed in this slot							
10	No hard disk installed in this slot							
11	No hard disk installed in this slot							
12	No hard disk installed in this slot							
13	No hard disk installed in this slot							
14	No hard disk installed in this slot							
15	No hard disk installed in this slot							

ENTER: Add/Remove Disk

Enclosure Slots								
Slot	Chn	ID	LUN	Vendor	Product	Attr.	Cap(MB)	Drive
0	SCSI-A 0	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 4
1	SCSI-B 15	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 5
2	SCSI-C 14	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 6
3	SCSI-D 0	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 7
4	SCSI-E 1	0	0	i	IBM DPSS-309170M	RW	8747	Drive 0
5	SCSI-E 3	0	0	i	IBM DPSS-309170M	RW	8747	Drive 1
6	SCSI-E 6	0	0	i	IBM DDYS-T18350M	RW	17501	Drive 2
7	No hard disk installed in this slot							

Select Disk								
Chn	ID	LUN	Vendor	Product	Attr.	Cap(MB)	Drive	
SCSI-F 6	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	Drive 3	

ENTER: Select Disk

12	No hard disk installed in this slot							
13	No hard disk installed in this slot							
14	No hard disk installed in this slot							
15	No hard disk installed in this slot							

ENTER: Add/Remove Disk

It should be outlined again, that SAF-TE and SES are definitely the better choices for auto hot pluggable subsystems. Only with these systems a secure Auto Hot Plug is guaranteed.

With the non-intelligent subsystems we highly recommend you to use only best quality components (disk shuttles, cables, terminators, etc.).

C.4.3.5 Menu Advanced Setup: Configure Controller, Advanced Settings

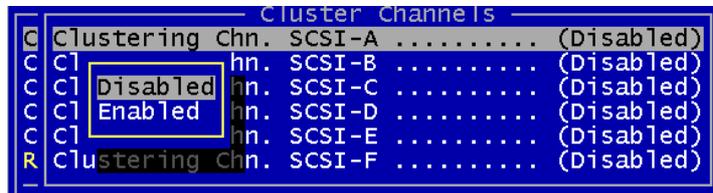
Within Advanced Settings there are three settings which control the configuration and address behavior of the ICP Controller's BIOS and DPMEM. Normally, there is nothing to change here.



Parameter	Alternatives	Description
Shrink BIOS after Post	Off	The BIOS is not shrunk after the post phase
	On	The BIOS is always shrunk after the post phase
	Auto	The ICP Controller decides on the BIOS shrink
BIOS RAM allocation method	Auto	The BIOS address space is automatically allocated
	Older	The BIOS address space is allocated according to an older PCI specification
DMEM mapping	Move below 1MB	The DPMEM address space is allocated below 1MB
	Do not move	The DPMEM address space is always above 1MB

C.4.3.6 Menu Advanced Setup: Configure Controller, Cluster Settings

This menu is selectable only with ICP Controllers which are equipped with Cluster RAIDYNE® and allows the user to enable the IO channels of the ICP Controller for clustering. Special care must be taken, that the IO-Processors of a shared channel (SCSI or Fibre Channel) have different IDs, otherwise there will be conflicts.



C.4.3.7 Menu Advanced Setup: Configure Controller, Clear Log Buffer

The ICP Controller records certain events in a logging buffer which is part of the Flash-RAM. If it is planned to use the ICP Controller in a completely new system, it is sensible to clear all events in this buffer. The Clear Log Buffer function detects possible entries and deletes them.

C.4.4 Menu Advanced Setup: Configure Physical Devices

After selection of this menu option ICPCON displays a list of the IO channels (SCSI or Fibre Channel) and the available devices. This list can be scrolled with the cursor up/down keys. The first column displays the IO channel ("SCSI" stands for SCSI and "FCAL" for Fibre-Channel). The next column shows the ID followed by the LUN (Logical Unit Number; normally always 0). The next column is either filled with a small "i" (for initialized), or empty. Initialization means the creation of ICP-specific configuration data on the hard disks and the adjustment of hard disk specific operating parameters (protocol, transfer rate, etc.). The next entry lists the vendor and type of device, followed by the drive's attributes (RW for

Select Physical Drive							
Chn	ID	LUN	Vendor	Product	Attr	Cap(MB)	Drive
SCSI-A	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	
SCSI-A	9	0	i	SEAGATE ST39102LC	RW	8682	
SCSI-A	10	0	i	FUJITSU MAF3364LC	RW	34746	
SCSI-A	11	0	i	FUJITSU MAF3364LC	RW	34746	Drive 4
SCSI-A	12	0	i	FUJITSU MAF3364LC	RW	34746	
SCSI-A	7	0		SCSI I/O Processor			
SCSI-B	8	0		SDR, Inc. GEM312 REV001			
SCSI-B	1	0	i	SEAGATE ST39175LC	RW	8682	
SCSI-B	2	0	i	SEAGATE ST39175LC	RW	8682	Drive 6
SCSI-B	3	0	i	SEAGATE ST39175LC	RW	8682	
SCSI-B	4	0	i	SEAGATE ST39175LC	RW	8682	
SCSI-B	15	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	
SCSI-B	7	0		SCSI I/O Processor			
SCSI-C	14	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	
SCSI-C	7	0		SCSI I/O Processor			
SCSI-D	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	
SCSI-D	7	0		SCSI I/O Processor			

[- F4: Extended Information, F5: Identify Drive, F8: Repair

Read/Write; RO for Read Only). The IO-Processors and SAF-TE-Processors (in our example channel B ID 8: SDR, Inc. GEM312 REV001) to not have attributes. The next column lists the capacity of the hard disk in MB (1MB = 1024KB; 1GB=1024MB). The last column gives information on the assignment of this Physical Device to a Logical/Array/Host Drive.

With ICP Fibre Channel RAID Controllers, F4 delivers extended information on Fibre Channel drives (World Wide Name, firmware-level and the Link- status of the drive).

If you press F5, the ICP Controller periodically turns the LED (if available) of the corresponding drive on and off. This helps to physically identify a specific drive in an array of many drives.

If you press F8 you can repair the configuration data of a physical drive. If your ICP Controller should ever display during his boot sequence a message like "Detected Primary Configuration Data error, using Secondary" or similar, you can try to repair the data with F8.

If you press ENTER on a specific drive ICPCON opens the "Configure Disk" menu. Here, all relevant adjustments and modifications concerning the physical drives can be carried out.

Note: If you create new Host Drives with the Express Setup function, ICPCON and the ICP Controller will automatically adjust all parameters to their respective best values. The transfer rate per drive depends on circumstances like cables, terminators, etc.. I.e., if a certain configuration does not allow 160 MB/sec, the drive will be automatically "throttled" down to the next trouble-free transfer rate.

Select Physical Drive							
Chn	ID	LUN	Vendor	Product	Attr	Cap(MB)	Drive
SCSI-A	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	
				Configure Disk	LC	8682	
				SCSI Parameter/Initialize	LC	34746	
				Format Disk	LC	34746	Drive 4
				Check Surface	LC	34746	
				View Defects/Status	LC		
				Deinitialize Disk	LC		
				Lock Disk	LC	8682	
				Unlock Disk	LC	8682	Drive 6
				Enclosure Status	LC	8682	
				Enclosure Slot Configuration	LC	8682	
SCSI-C	7	0		SCSI I/O Processor			
SCSI-D	0	0	i	QUANTUM ATLAS10K2-TY184J	RW	17509	
SCSI-D	7	0		SCSI I/O Processor			

[- F4: Extended Information, F5: Identify Drive, F8: Repair

C.4.4.1 Menu Advanced Setup: Configure Phys. Devs., SCSI Parameter /Initialize

This option can destroy all data on the hard disk.

If a hard disk is not yet initialized, you have to initialize it first. ICPCON copies ICP specific configuration blocks on the hard disk, a primary block and a mirrored secondary block. The possible settings are different if you select a SCSI hard disk or a Fibre Channel hard disk. With a FCAL

Chn	ID	LUN	Vendor	Product	Attr	Cap(MB)	Drive
SCSI-A	0	0	QUANTUM	ATLAS10K2-TY184J	RW	17509	
				Configure Disk	LC	8682	
				SCSI Parameter/Initialize	LC	34746	
Fo				Initialize Disk		34746	Drive 4
Ch				Sync. Transfer	(Enabled)	34746	
Vi				Sync. Transfer Rate	(160.0 MB/s)		
De				Disconnect	(Enabled)		
				Disk Read Cache	(On)	8682	
Lo				Disk Write Cache	(On)	8682	Drive 6
Un				Tagged Queues	(On)	8682	
				F4: Advanced Configuration		8682	
				Enclosure Status	K2-TY184J RW	17509	
				Enclosure Slot Configuration	sor		
SCSI-C	7	0		SCSI I/O Processor			
SCSI-D	0	0	QUANTUM	ATLAS10K2-TY184J	RW	17509	
SCSI-D	7	0		SCSI I/O Processor			

hard disk there are only a few settings which are relevant. You should always check that they are all "Enabled" or "On". If the hard disk you have selected is an Ultra 160 device and the ICP Controller has Ultra 160 SCSI channels, F4 "Advanced Configuration", allows you to configure the so-called "Domain Validation" (a cyclical check of the correct data transfer at a given rate). Normally this parameter is On.

The various settings have the following meanings:

Sync. Transfer: Enabled

The SCSI-bus knows two methods of data transfer: asynchronous and synchronous transfer. Each SCSI device must be able to perform the first type of transfer, the second one is optional. The advantage of the synchronous transfer consists in a higher data transfer rate, since the signal transfer times on the possibly long SCSI-cable have no influence on the transfer rate anymore. Two SCSI-bus participants which want to exchange data between each other have to check if and how (i.e., with which parameters) a synchronous data transfer between them is possible. Therefore, the mere setting does not automatically enable synchronous data transfer; this mode is only effective if both devices support it and after they have checked their capability of communicating with each other in this mode.

Sync. Transfer Rate

This is the synchronous data transfer rate in MB/s. Ultra 160 SCSI allows up to 160 MB/s. If a given SCSI-cable does not allow a certain transfer rate, it can be reduced to a value that allows a trouble-free data transfer. The reason for such a restriction is not necessarily a "bad" SCSI-cable. Lowering the transfer rate may also become necessary when you set up a special configuration which does not allow full speed.

Disconnect: Enabled

The concept of the SCSI-bus allows several participants. It is particularly important to guarantee a high degree of action overlapping on the SCSI-bus. This high degree of overlapping becomes possible when a SCSI device has an enabled "Disconnect".

Disk Read Cache / Disk Write Cache / Tagged Queues

Out of performance reasons, the read ahead and write cache of the hard disk should be always On. Tagged Queues is a SCSI feature which allows the drive to execute more than one command at a time, if available, it should also be On.

If you leave this configuration form with <ESC> and you have made changes, ICPCON displays a security request. The warning of the destruction of all data implies different evaluations, depending on the device's current state and the options you selected:

```
Initialization of disk will destroy all data.
Continue ? (Y/N)
```

First Initialization of the Device. In this case, the warning must be taken seriously. If the drive was previously connected to a different controller (e.g. NCR etc.) and still contains important data, this data will be lost now.

The Device was already initialized. If only internal parameters such as Disconnect, Synchronous Transfer, and Disk Read/Write caches, or tagged queues have been changed, the data on the drive remains intact. Only the function state of the device changes.

C.4.4.2 Menu Advanced Setup: Configure Phys. Devs., Format Disk

This option destroys all data on the hard disk.

All manufacturers of hard disks deliver their products already formatted and surface-tested. For new hard disks it is neither necessary, nor advisable to perform the Format Disk. This procedure is only indicated if you have doubts on the hard disk's condition.

The time required for the Format Disk of a hard disk depends on the hard disk itself. It can take quite a long time (up to days !). Often it seems that nothing happens and that the system hangs (no LED indication). If you put your ear on the hard disk you can hear the actuator stepping (with some drives one step per minute or longer). **Never interrupt a Format Disk procedure.** This may lead with a very high probability to a non-functioning hard disk. Before the actual formatting, ICPCON asks you whether the "Grown Defect" table of the hard disk should be deleted. Some users believe that this makes a hard disk with a lot of grown defects like new. This is wrong. As soon as the bad sectors are accessed again, a re-assign will happen, generating a new grown defect.

C.4.4.3 Menu Advanced Setup: Configure Phys. Devs., Check Surface

This option destroys all data on the hard disk.

This option allows the checking of the surfaces of the hard disk media. The ICP Controller writes and reads certain data patterns and checks them for correctness.

After confirming the security request, a progress information is displayed. You can interrupt the Check Surface option by pressing <ESC>.

C.4.4.4 Menu Advanced Setup: Configure Phys. Devs., View Status/Defects

Grown defects. Number of media defects that have occurred in addition to the media defects the hard disk already had upon delivery.

Primary defects. Number of media defects that the hard disk already had upon delivery.

Last status: The Last Status gives detailed information on the last failure of a hard disk. The information is only present until the next hard reset of the system and may help for deeper failure analysis or tracing.

```
Grown defects: 0
Primary defects: 1358
Last status: 0x00000000
```

C.4.4.5 Menu Advanced Setup: Configure Phys. Devs., Deinitialize Disk

This option destroys all data on the hard disk.

This menu option allows you to de-initialize a hard disk which has previously been initialized for use with the ICP Controller. By doing so, the specific ICP information present on the device is removed. Obviously, the de-initialization cannot restore data that was lost during initialization.

C.4.4.6 Menu Advanced Setup: Configure Phys. Devs., Lock/Unlock Disk

This option is only high-lighted when you have selected a removable hard disk (e.g., Syquest, Iomega). Before you can initialize a cartridge you have to lock it. Before removing it you have to unlock it.

C.4.4.7 Menu Advanced Setup: Configure Phys. Devs., Enclosure Status

After selecting this option you can either view the Enclosure Status or view/configure the enclosure slots. Before you can use the Auto Hot Plug with a SAF-TE subsystem, you first have to configure the subsystem (more precisely it's intelligence, the so-called SEP - SAF-TE Enclosure Processor).

Enclosure Status	
Audible Alarm Status	(Not available)
Door Lock Status	(Not available)
Power supply 0 status	(Operational and on)
Power supply 1 status	(Operational and on)
Fan 0 status	(Operational)
Fan 1 status	(Operational)
Fan 2 status	(Operational)
Enclosure temperature	(OK)
Temperature sensor 0	(33 °C / 91 °F OK)

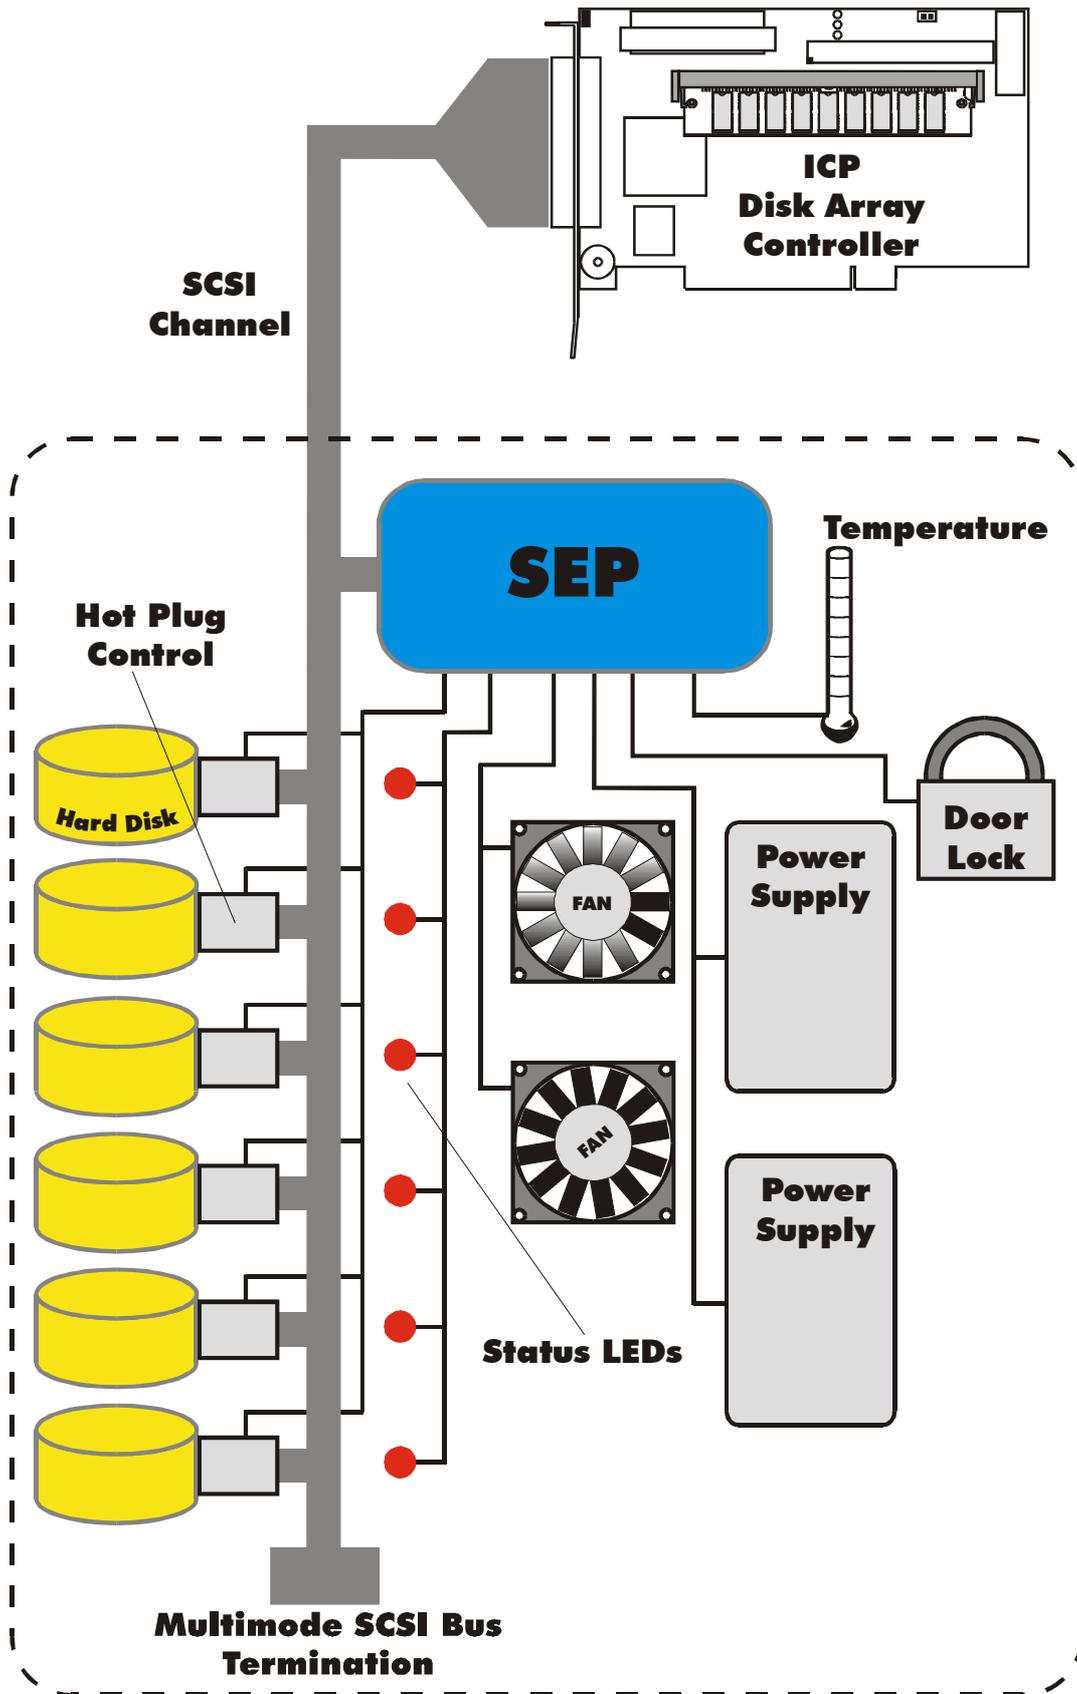
F10: Refresh Info

To assign a specific hard disk to a SAF-TE enclosure slot, press ENTER on the empty slot and choose one of the hard disks.

Enclosure Slots								
Slot	Chn	ID	LUN	Vendor	Product	Attr.	Cap(MB)	Drive
0	SCSI-A 9	0	i	SEAGATE	ST39102LC	RW	8682	
1	SCSI-A 10	0	i	FUJITSU	MAF3364LC	RW	34746	
2	SCSI-A 11	0	i	FUJITSU	MAF3364LC	RW	34746	Drive 4
3	SCSI-A 12	0	i	FUJITSU	MAF3364LC	RW	34746	
4	SCSI-B 1	0	i	SEAGATE	ST39175LC	RW	8682	
5	SCSI-B 2	0	i	SEAGATE	ST39175LC	RW	8682	Drive 6
6	SCSI-B 3	0	i	SEAGATE	ST39175LC	RW	8682	
7	SCSI-B 4	0	i	SEAGATE	ST39175LC	RW	8682	

ENTER: Add/Remove Disk

The next page shows a block diagram of a SAF-TE subsystem.



C.4.5 Menu Advanced Setup: Configure Logical Drives

Logical Drives (hierarchy level 2) are installed in this main menu option.

Selecting *Configure Logical Drives* leads you to the screen shown next.

The F4 key gives you a list of all the hard disks this Logical Drive consists of. If it is a Logical Drive of the type **Disk**, it only consists of one single hard disk. If a Logical Drive consists of more hard disks, it is of the type **Chain** (concatenation of several hard disks).

F3 loads the media of removable hard disks and F5 turns for some seconds periodically the hard disks LED on and off.



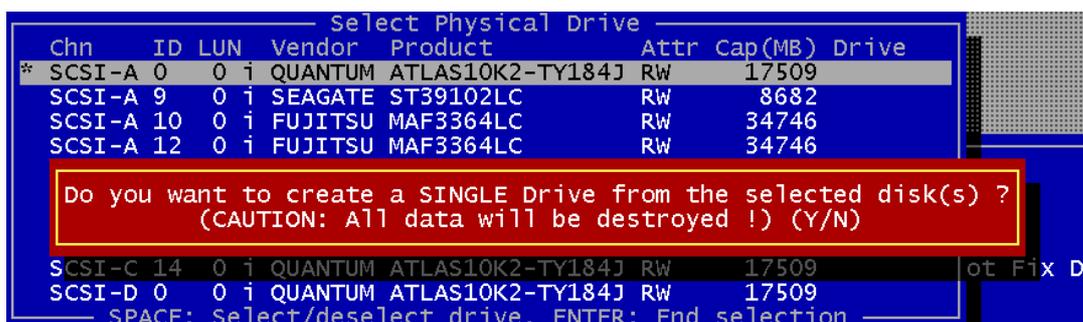
To configure and setup an new Logical Drive, select Create new Logical Drive. ICPCON displays a list of free Physical Drives (which are not already part of a Logical drive).

Select Physical Drive

Chn	ID	LUN	Vendor	Product	Attr	Cap(MB)	Drive
SCSI-A 0	0	i	QUANTUM	ATLAS10K2-TY184J	RW	17509	
SCSI-A 9	0	i	SEAGATE	ST39102LC	RW	8682	
SCSI-A 10	0	i	FUJITSU	MAF3364LC	RW	34746	
SCSI-A 12	0	i	FUJITSU	MAF3364LC	RW	34746	
SCSI-B 1	0	i	SEAGATE	ST39175LC	RW	8682	
SCSI-B 3	0	i	SEAGATE	ST39175LC	RW	8682	
SCSI-B 4	0	i	SEAGATE	ST39175LC	RW	8682	
SCSI-B 15	0	i	QUANTUM	ATLAS10K2-TY184J	RW	17509	
SCSI-C 14	0	i	QUANTUM	ATLAS10K2-TY184J	RW	17509	
SCSI-D 0	0	i	QUANTUM	ATLAS10K2-TY184J	RW	17509	

SPACE: Select/deselect drive, ENTER: End selection

If you select one Physical Drive with SPACE and press ENTER, ICPCON suggests to create a SINGLE Drive out of this hard disk.



If you confirm with <Y>, ICPCON allows you to limit the size of the Logical Drive. This becomes interesting when you configure later on an Array Drive with several identical Logical Drives and you want to make sure that you get appropriate spare hard disks in the future. It would be bad luck if the new hard disk would have 17508MB, only. It simply wouldn't fit into the Array Drive. If you limit the capacity to e.g., 17000MB from the beginning, you can be sure that all future 18GB hard disk will have at least 17000MB and thus can be used as spare hard disk.



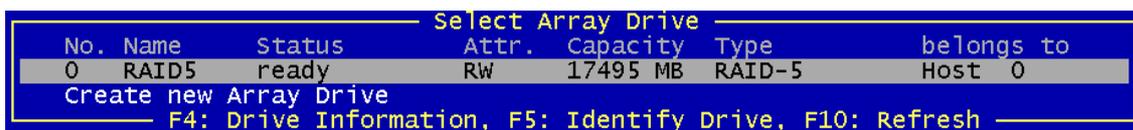
If you select two or more Physical Drives with SPACE, ICPCON suggest the creation of a Logical Drive of the type *Chain*. In some literature Disk Chaining is also called *Disk Spanning*. You can picture the functioning mechanism of a type *Chain* Logical Drive as follows: all hard disks forming the Logical Drive are linked together one by one in the exact same order in which they have been selected with the SPACE bar. This concatenation can be compared with a chain. If, for example, the Logical Drive consists of 4 hard disks with 2000MB each, the Logical Drive will have a capacity of 8000MB. When data is written to this Logical Drive, the first hard disk is filled first, then the second, and so on. Although it is not advisable, Logical Drives of the type Chain, can also be components of Array Drives.

C.4.6 Menu Advanced Setup: Configure Array Drives

This main menu option allows you to configure Array Drives (level of hierarchy 3). Array Drives with the following listed RAID levels can be configured within this menu.

- RAID 0 pure data striping without redundancy
- RAID 1 disk mirroring
- RAID 4 data striping with dedicated parity drive
- RAID 5 data striping with striped parity
- RAID 10 RAID 0 combined with RAID 1

The ICP Controller can manage up to 35 Array Drives (with different RAID levels) simultaneously. Obviously, the physically existing number of hard disks will limit the number of parallel used Arrays.

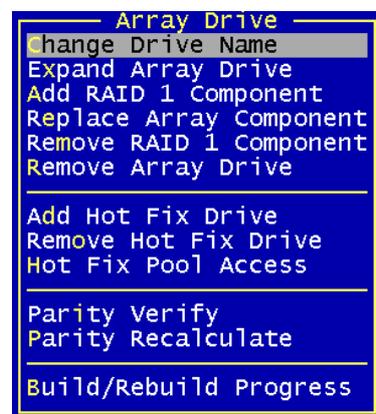


No.	Name	Status	Attr.	Capacity	Type	belongs to
0	RAID5	ready	RW	17495 MB	RAID-5	Host 0

Create new Array Drive
F4: Drive Information, F5: Identify Drive, F10: Refresh

F4 displays level by level detailed information on the selected Array Drive (the structure, the order, which hard disks are part of the Array Drive). With F5 the ICP Controller turns for a few seconds the LEDs of all hard disks belonging to this Array Drive periodically on and off. This may help to identify the hard disks.

If you press ENTER on a specific Array Drive ICPCON displays a list of possible options.



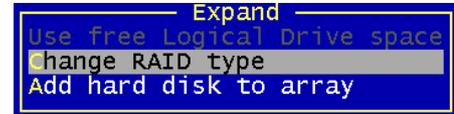
change Drive Name
Expand Array Drive
Add RAID 1 Component
Replace Array Component
Remove RAID 1 Component
Remove Array Drive
Add Hot Fix Drive
Remove Hot Fix Drive
Hot Fix Pool Access
Parity Verify
Parity Recalculate
Build/Rebuild Progress

C.4.6.1 Menu Advanced Setup: Configure Array Drives, Change Drive Name

This command allows you to change the name of an Array Drive. The name serves to identify an Array Drive in ICPCON or ICP RAID Navigator. This can be very helpful for configurations where several Host Drives of various types are operated by a single controller.

C.4.6.2 Menu Advanced Setup: Configure Array Drives, Expand Array Drive

The *Expand Array Drive* option which is also available with ICP RAID Navigator, contains three major functions:



1. *Conversion* of a free space on the Logical Drives into a separate Host Drive
2. *Migration* of the RAID level of a given Array Drive
RAID 0 -> RAID 4 and vice versa
RAID 0 -> RAID 5 and vice versa
3. *Expansion* of the capacity of a given Array 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.

C.4.6.3 Menu Advanced Setup: Configure Array Drives, Add RAID 1 Component

In certain "emergency" cases this is a very powerful and helpful option. This function allows you to add to a Logical Drive which is member of an Array Drive, another Logical Drive as a mirror drive (RAID-1).

Example: You have configured an Array Drive with 4 Logical Drives. One Logical Drive has failed and the Array Drive went into the fail state. Another failure would cause data loss. Unfortunately, you find another Logical Drive, which is shortly before failing (e.g., you hear a strange noise from it, or it's grown defect counter explodes). If you now initiate a hot plug it is very likely that this critical Logical Drive will also fail. This would result in a disaster. To avoid that problem, you can mirror in a first step a new good Logical Drive to the critical one. When the copying is finished you remove the critical Logical Drive and then carry out a hot plug procedure.

C.4.6.4 Menu Advanced Setup: Configure Array Drives, Replace Array Component

If a Logical Drive of an Array Drive without a Hot Fix drive should fail (or is very likely to fail, soon), you should replace the defective hard disk with a new one as soon as possible because the Array Drive is without redundancy. The replacement Logical Drive has to have at least the same capacity as the failed one. The replacement is carried out either with ICPCON or ICP RAID Navigator. After having installed the replacement hard disk as a new Logical Drive, you can add it to the Array Drive. After selecting the Logical Drive which needs to be exchanged, ICPCON offers a list of existing Logical Drives which can be used as replacement units. The Array Drive's state is changing into *rebuild* and the missing data is automatically reconstructed on the new Logical Drive.

C.4.6.5 Menu Advanced Setup: Configure Array Drives, Remove RAID 1 Component

This option corresponds with the *Add RAID-1 Component* option. It allows you to remove a previously configured RAID-1 combination.

C.4.6.6 Menu Advanced Setup: Configure Array Drives, Remove Array Drive

This command allows you to remove an existing Array Drive.

All the data of the Array Drive will be lost !

Before you confirm the security request with <Y>, you should be sure about this choice.

Do you really want to remove the selected drive ? (Y/N)

Note: If an Array Drive has been removed, it can perhaps be rebuilt without data loss if it is reconstructed in the exact same order it had been built before, if the components of the Array Drive, that is the Host Drives, have not been modified, if the stripe size and RAID level is the same and if a Non-Destructive Build is carried out. In all other cases ALL DATA WILL BE LOST !

C.4.6.7 Menu Advanced Setup: Configure Array Drives, Add Hot Fix Drive

This submenu option allows you to add a Hot Fix drive to an existing RAID 1, RAID 4, RAID 5, or RAID 10 Array Drive. There are two different types of Hot Fix drives: *Private* and *Pool* Hot Fix drives. 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 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. A *Private* Hot Fix drive is dedicated to one RAID 1, RAID 4, RAID 5 or RAID 10 Array Drive.

Only drives that meet the following requirements are suitable as Hot Fix drives:

1. The Logical Drive that is to become a Hot Fix drive must not be an active component of another Array Drive.
2. The Logical Drive that is to become a Hot Fix drive must have a storage capacity greater than or equal to the storage capacity of the smallest Logical Drive of the Array Drive. Example: A type RAID 5 Array Drive consists of the following components:

Logical Drive 0	2000MB
Logical Drive 1	1500MB
Logical Drive 2	1100MB
Logical Drive 3	2000MB

This Array Drive has a usable storage capacity of 3300MB. A Hot Fix drive for this array must have at least 1100MB of storage capacity. (Note: in order not to waste valuable storage capacity, it is strongly recommended that all Logical Drives forming an Array Drive have the same storage capacity.)

What happens after a drive failure ?

The controller will substitute a failed Logical Drive with a Hot Fix drive only if the Array Drive was in the *ready* state before the failure, or, in other words, a Hot Fix drive can only be activated if the corresponding Array Drive had a state of data redundancy at the moment of failure.

1. After a short while, the controller's alarm turns on.
(Note: the alarm is activated only when the Array Drive is being accessed.)
2. The controller activates the *fail* operation mode. In this mode, the Array Drive remains fully operational. The data located on the failed drive is generated by means of the redundancy information stored on the other drives, without causing any decrease in performance.
3. The controller integrates the Hot Fix drive into the Array Drive and starts to reconstruct the data and redundancy information. The Array Drive is now in the *rebuild* operation mode.

Obviously, no other hard disk may fail until all data has entirely been reconstructed on the Hot Fix drive, because up to that moment, the system is operating without redundancy.

Notes: In some literature, Hot Fix drives are also called *Hot-Spare* drives. You can add or remove Hot Fix drives also with the ICP RAID Navigator.

C.4.6.8 Menu Advanced Setup: Configure Array Drives, Remove Hot Fix Drive

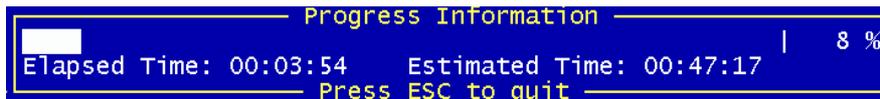
This option allows you to remove a Hot Fix Drive from an existing Array Drive. Naturally, the Hot Fix drive must not be used up so far.

C.4.6.9 Menu Advanced Setup: Configure Array Drives, Hot Fix Pool Access

By selecting the *Hot Fix Pool Access* option, the access of a specific Array Drive to the Hot Fix pool can be enabled or disabled.

C.4.6.10 Menu Advanced Setup: Configure Array Drives, Parity Verify

The redundancy information which is calculated during an array *build* or *rebuild* is stored on a dedicated Logical Drive (RAID 4), or is distributed over all Logical Drives of the Array Drive (RAID 5). This information is often called *parity data*. The calculation is made with an exclusive OR function (XOR). If a Logical Drive of an Array Drive fails, its data can be recalculated by means of the data present on the other Logical Drives of the Array Drive and the parity data. The *Parity Verify* function allows you to check the consistency of an Array Drive's parity data. The verification may take quite a long time, but you can terminate it by pressing <ESC>.



C.4.6.11 Menu Advanced Setup: Configure Array Drives, Parity Recalculate

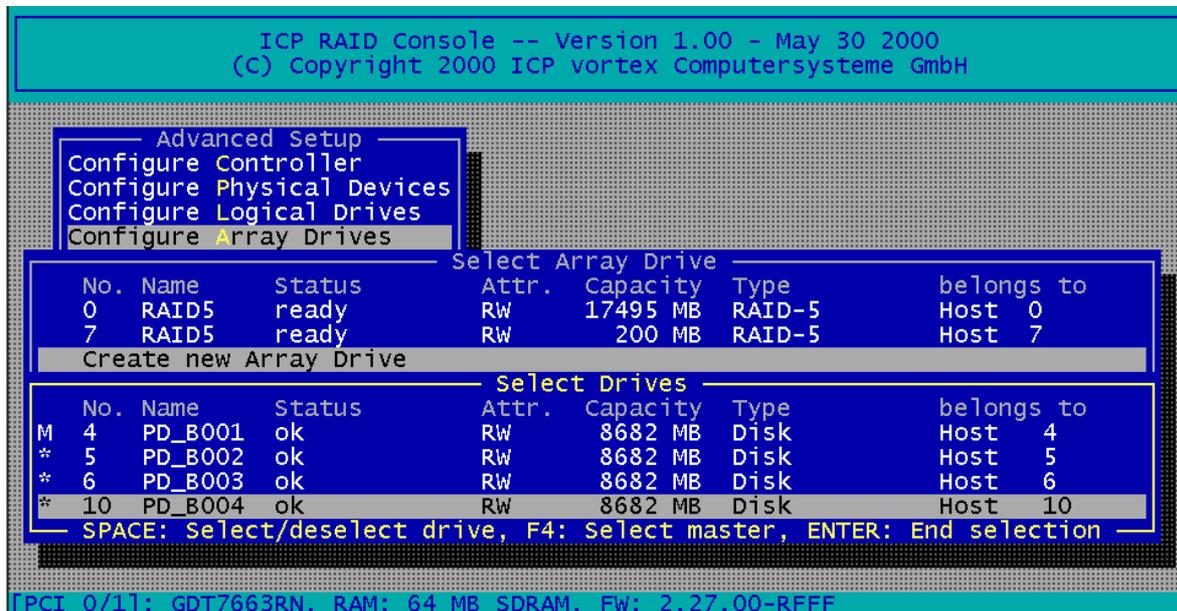
If the parity information of a given Array Drive is defective, this function may be used to recalculate it anew.

C.4.6.12 Menu Advanced Setup: Configure Array Drives, Build/Rebuild Progress

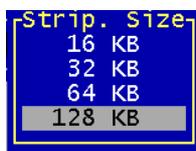
Whenever an Array Drive is in the *build* or *rebuild* state, you can select this option, to get progress information and estimates for the required time.

C.4.6.13 Menu Advanced Setup: Configure Array Drives, Create new Array Drive

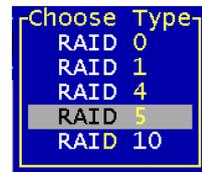
After pressing ENTER, ICPCON lists all free Logical Drives, which are free (not yet part of Array / Host Drives). The selection bar can be moved with the cursor up/down keys and the Logical Drives can be selected/deselected with the SPACE bar.



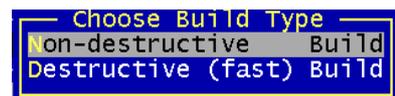
The "M" means Master. For a striping array (RAID 0, 4, 5, 10), this is the first Logical Drive in the array. For a RAID 1 (mirroring) array this is the Logical Drive which contains the valid data and which should be copied to the second Logical Drive. After pressing ENTER



ICPCON displays a list of possible RAID levels. ICPCON will ask you for the Stripe Size. This is the size of the stripes into which the data is divided. Valid values are 16KB, 32KB, 64KB or 128KB. The default is 128KB.



The Build Type decides how the redundancy information is created. The Destructive Build (which is only available if ICPCON was loaded with <CTRL><G>) fills all hard disks with "0" and no further parity calculation is necessary. This build method is pretty fast, but the build must be successfully finished before you can leave ICPCON (otherwise the build mode is automatically changed to Non-destructive during the next boot). The Non-destructive build really calculates the redundancy information from the contents of the data stripes. It takes longer, but can run in the background, i.e., during the installation of the operating system.



This security request has to be taken seriously. If you confirm with <Y> all data are lost.

Do you want to create a Array Drive from the selected drives ?
(CAUTION: All data will be destroyed !) (Y/N)

If necessary you can limit the Array Drive's capacity. For this example we take the complete capacity.

Drive Size (1..26047 MB): 26047

The Array Drive has entered the *build* state, i.e., the parity information is currently generated. After completion of the build process, the Array Drive's state is *ready*, i.e., fault tolerant.

Select Array Drive						
No.	Name	Status	Attr.	Capacity	Type	belongs to
0	RAID5	ready	RW	17495 MB	RAID-5	Host 0
4	RAID5	build	RW	26047 MB	RAID-5	Host 4
7	RAID5	ready	RW	200 MB	RAID-5	Host 7
Create new Array Drive						
F4: Drive Information, F5: Identify Drive, F10: Refresh						

Notes on the Configuration of RAID 0, 1, 4, 5 and 10 Arrays Drives

(1) Use preferably Logical Drives of the type *disk* to build an Array Drive.

Of course, RAID Array Drives can be configured with Logical Drives of the type *chain*, too, but the aspects of security should be taken into consideration as well. For regular RAID Array Drives, type *disk* Logical Drives are used.

(2) The Logical Drives of an Array Drive should have the same storage capacity.

In order not to waste valuable storage capacity, you should only use Logical Drives that have the same storage capacity for an Array Drive.

(3) The Hot Fix drive provides the utmost security.

One of the reasons for which RAID Array Drives are used definitely lies with the redundancy they provide, that is, the data security you still have even in the event of a hard disk failure, thus resting assured against loss of data and time. For the purpose of the following considerations, we define the term ***time without redundancy, TWR***. Set apart the time needed to set up the Array Drive (state *build*), the time without redundancy should be kept as short as possible. Let us assume that one of the hard disks of a RAID 5 Array Drive fails. The Array Drive is without redundancy. TWR starts to run. Any superfluous prolongation of the TWR (because you have to get a replacement hard disk, or because you did not realize the failure immediately since you didn't hear the ICP Controller's alarm signal, or because nobody checked the file server) increases the risk of data loss which will occur if a second hard disk should fail. Therefore, new redundancy should be created as soon as possible and in an entirely automated manner. Integrating a Hot Fix drive as an immediately available and auto-replacing hard disk is the only way to keep the TWR as short as possible. Only a Hot Fix drive can ensure optimal Array Drive security and constant data availability. Of course a Hot Fix drive is not mandatory. If you control the Array Drive at regular intervals and immediately replace a defective hard disk (by shutting down the system or Hot Plug), you can do without a Hot Fix drive.

C.4.7 Menu Advanced Setup: Configure Host Drives

This main menu option allows you to configure Host Drives (level of hierarchy 4). As already mentioned before, these are the drives the Host Computer is aware of. Host Drives can consist of a single hard disk, or of many hard disk combined to a RAID 5 Array Drive. The menus Split/Merge Host Drives and Partition Host Drives are only available if ICPCON was loaded with <CTRL><G> from the ICP Controller's Flash-RAM. The additional capacity resulting from an Online Capacity Expansion is shown as another Host Drive. If you expanded the capacity of the Array Drive a second time, there would be three Host Drives, belonging to one and the same Array Drive. Since there is currently no operating system, which supports "growing hard disks", this expansion method is the only safe way to introduce new capacity.

Host Drive	
Change Drive Name	
Swap Host Drives	
Remove Host Drive	

Split Host Drive	
Merge Host Drives	

Partition Drive	
Overwr. Master Boot Code	

Drive Type (cluster)	

C.4.7.1 Menu Advanced Setup: Configure Host Drives, Change Drive Name

This command allows you to change the name of a Host Drive. The name serves to identify a Host Drive within ICPCON and ICP RAID Navigator.

C.4.7.2 Menu Advanced Setup: Configure Host Drives, Swap Host Drives

When the PCI computer is switched on, the Host Drives are initialized in the order of the Host Drive list, which means that the operating system is booted from the Host Drive having the lowest number. For reasons of flexibility, a Host Drive's position in the list can be changed. However, the position of the Host Drive from which the operating system is booted and the position of the Host Drive from which ICPCON (disk version) was started (both can be the same), cannot be changed. If you wish to change the position of these drives, you have to boot the operating system and ICPCON from a floppy disk or use the ICPCON loadable from the Flash-RAM of the controller. To change the position of a Host Drive in the Host Drive list, highlight the Host Drive and confirm with ENTER. Then, type on the new position and press ENTER.

C.4.7.3 Menu Advanced Setup: Configure Host Drives, Remove Host Drives

Removing a Host Drive is a serious action. All data will be lost after removal.

If you want to remove a Host Drive belonging to an Array Drive for which several Host Drives exist (after capacity expansion, or after splitting), ALL other Host Drives will also be removed.

C.4.7.4 Menu Advanced Setup: Configure Host Drives, Split Host Drive

For some purposes it might be of interest to split an existing Host Drive into two or several Host Drives. Each Host Drive looks to the operating system just like a single hard disk. Since the new Host Drives have smaller capacities ICPCON has to write new header information on the two Host Drives. **All data will be lost.**

C.4.7.5 Menu Advanced Setup: Configure Host Drives, Merge Host Drives

This function reverses the *Split Host Drive* option. Only such Host Drives can be merged which belong to the same Array Drive or Logical Drive. Since the new Host Drive has a larger capacity ICPCON has to write a new header information on the new Host Drives. **All data will be lost.**

C.4.7.6 Menu Advanced Setup: Configure Host Drives, Partition Host Drives

The partitioning menu has similar functions as the MS-DOS program FDISK. You can create and delete a partition and also change the active partition. MS-DOS can only be booted from an active partition. Just like FDISK, ICPCON can handle primary partitions, extended partitions, and logical drives within the extended partitions.

C.4.7.7 Menu Advanced Setup: Configure Host Drives, Overwrite Master Boot Code

This option creates a valid and consistent master boot record on the selected Host Drive and should be carried out on any new Host Drive on which Windows NT is installed. **Never use it on Host Drives containing valid data. They will be lost.**

C.4.7.8 Menu Advanced Setup: Configure Host Drives, Drive Type (Cluster)

This option allows you to enable a Host Drive for clustering. Before that the shared IO channel(s) of the ICP Controllers have to be also enabled for clustering.

Chapter D

ICP RAID Navigator

D. ICP RAID Navigator

D.1 Introduction

The ICP RAID Navigator (ICPRNAV) is a powerful tool for setting up, monitoring and maintaining mass storage subsystems based on ICP Controllers. Different to ICP RAID Console the ICP RAID Navigator is a pure GUI-style application, designed for the operation under Windows 95, 98, NT and Windows 2000. The main features are:

Setup of hard disks controlled by the ICP Controller

- Setup and initialize Physical Drives, change the SCSI parameters and cache settings.
- Easy setup and installation of single disks or Array Drives consisting of one or more Physical Drives (chaining, RAID 0, 1, 4, 5 or 10).
- Online installation of Host Drives based on a single disk or an Array Drive.

Maintenance and tuning of existing Array Drives, Host Drives and the ICP Controller

- Configure the cache memory.
- Update the controller firmware.
- Change the device parameters.
- Check or recalculate the parity data of RAID 4 and RAID 5 Array Drives.
- Online capacity expansion of Array Drives.
- Split and merge Host Drives.
- Hot Plug. Replace a failed member of an Array Drive.
- Hot Fix. Configure a spare disk for an Array Drive.
- Configure SAF-TE compliant enclosures.

Monitoring of the whole subsystem

- Watch the performance and throughput of virtually every part of the subsystem. Simply drag-and-drop the device on the statistics window.
- Check the hard disks for retries and reassigns (replacements of defective sectors on a hard disk) to trace problems which might become potential in the future.
- Check the grown defect table of the hard disks to replace a defective drive in time.

Remote configuration and maintenance

- Except the Hot Plug function, all of the above features can be accessed via network. You can maintain and monitor a customer's server via the Internet.
- Supported protocols:
 - IPX/SPX
 - TCP/IP
 - NetBIOS
- The access to the ICP Controller can be protected by password. The password is encrypted.

To install the ICP RAID Navigator, use the icptools program on the ICP System CDROM.

D.2 The ICP RAID Navigator "Controls"

D.2.1 The Toolbar

The toolbar can be made visible or hidden by selecting "Toolbar" from the "View" menu:

By clicking on the different buttons you can open and close the windows of the corresponding programs of the ICP RAID Navigator:



The toolbar can be moved away from the top of the RAID Navigator window and is then shown in a small extra window. To place the toolbar back under the menu bar you can double click on the top window bar of the toolbar or drag & drop the toolbar back under the menu bar. You can also place the toolbar on the bottom of the RAID Navigator window.

Icon	Opens / closes	Description
	Select Controller	Select a local or remote ICP Controller for further actions.
	Physical Configuration	Show and/or modify the ICP Controller and device settings.
	Logical Configuration	Show and/or modify the Logical Drive configuration.
	Statistics	Show statistics.
	Events	Show the ICP Controller events.
		Show information on the ICP RAID Navigator.
		Click on this icon and then on the icon you want on-line help.

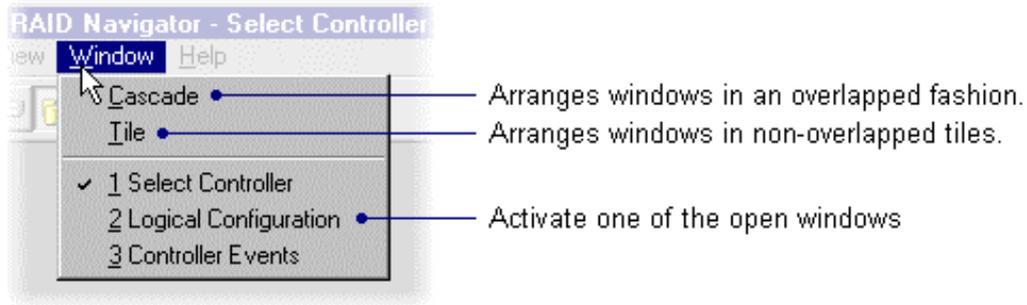
D.2.2 The Status Bar

The status bar at the bottom of the main window displays status information on the icon or the menu option the mouse pointer is currently placed on. The status bar can be displayed or hidden by selecting Status Bar from the View menu.



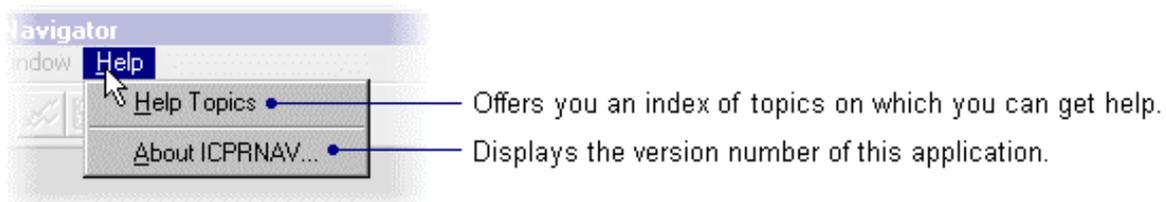
D.2.3 "Window" Menu Commands

These commands allow you to arrange the windows in the ICP RAID Navigator application window or to activate an open window.



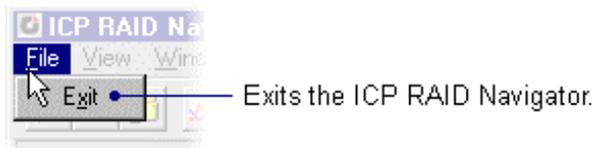
D.2.4 "Help" Menu Commands

The Help menu offers the following commands to provide you with online help:



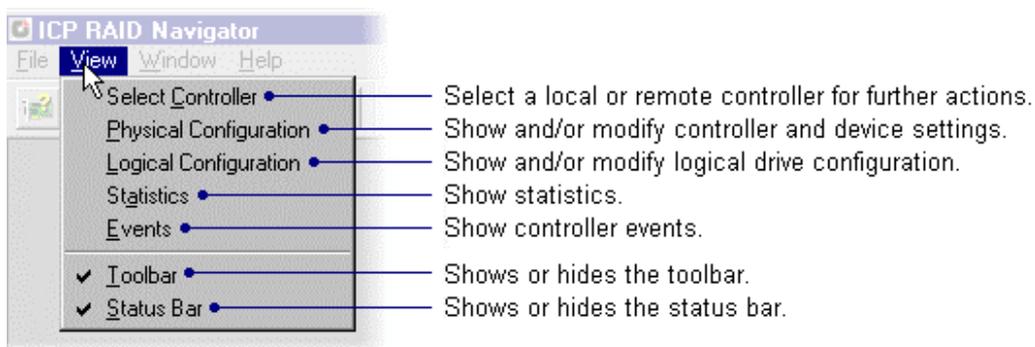
D.2.5 "File" Menu Commands

Here you can end your ICP RAID Navigator session. Shortcuts: Press Alt+F4 or click to close the window.



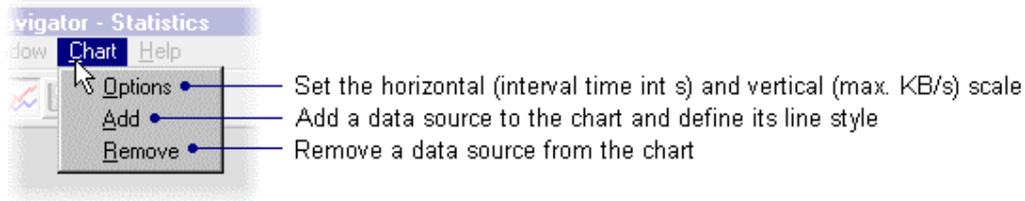
D.2.6 "View" Menu Commands

Use the items in this menu to open or close the windows of the main components of the ICP RAID Navigator or change the appearance of the main window.



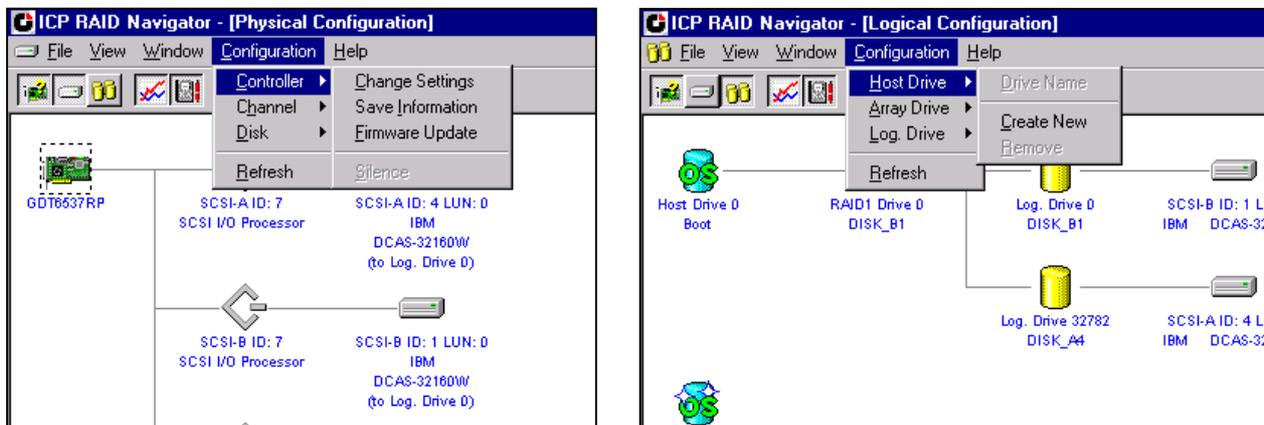
D.2.7 The "Chart" Menu

The chart menu appears when you open the statistics window. Here you can add and remove data sources from the chart and configure the chart.



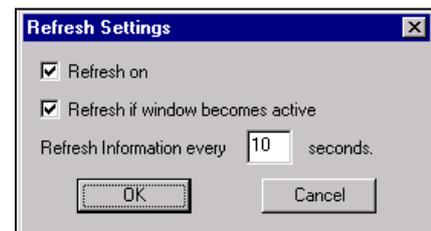
D.2.8 The "Configuration" Menu Commands

This menu appears if you have selected the Physical Configuration window or the Logical Configuration window.



Here you can set the refresh rate for the Physical or the Logical Configuration windows. This is the rate, which is used by the ICP RAID Navigator to update the contents of the physical and logical configuration windows.

The options are:



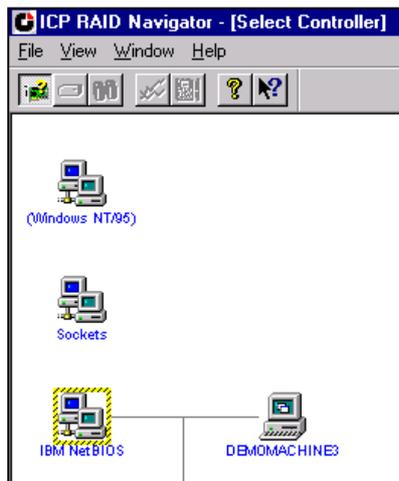
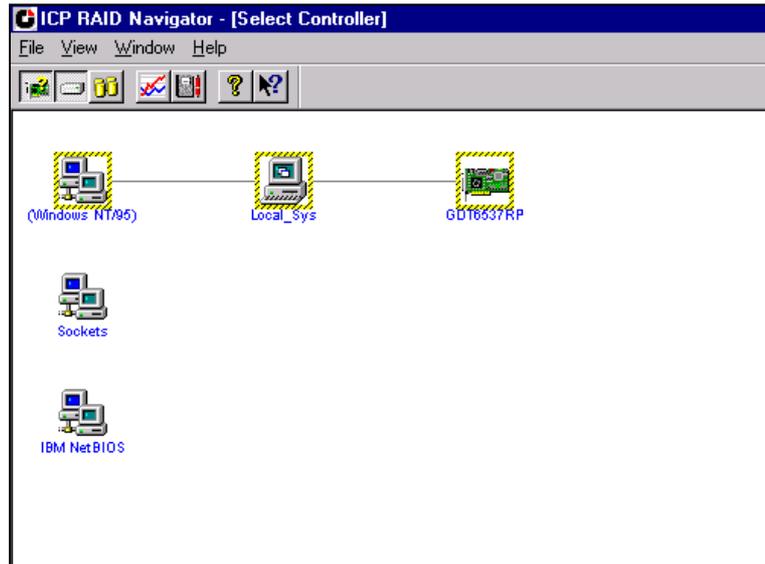
Refresh on	Activates and deactivates the window update.
Refresh if window becomes active	Update the window if it becomes active.
Refresh information every ... seconds	Sets the frequency for automatic window refreshes.

The menu options you can select here are also available if you click the right mouse button on the corresponding icon in the Physical or the Logical Configuration windows.

D.3 Select Controller

This window is used to select an ICP Controller for all further actions within the ICP RAID Navigator.

In a first step you should select the desired protocol (Windows NT/95/98, Sockets, IBM NetBIOS) for the communication between the ICP RAID Navigator and the system which is equipped with the ICP Controller by double clicking on the protocol icon on the very left side of the window. If you want so select an ICP Controller via TCP/IP (sockets), type in the corresponding TCP/IP address or the server system's name. You are then asked to enter a user name and a password. This information is transmitted with encryption.



After the protocol was selected, select the system which contains the ICP Controller and afterwards the ICP Controller itself. The selection and de-selection is done by double clicking on the ICP Controller. The selected protocol, system and ICP Controller are shown in a dashed yellow frame.

No more action can be done in this menu.



D.4 Physical Configuration Window

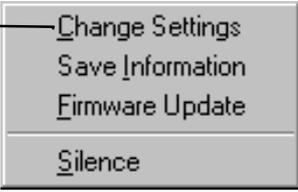


This window shows the physical configuration of the SCSI or Fibre Channel busses of the selected ICP Controller and the devices connected with these busses. The complete configuration is shown as a tree starting from the left with the ICP Controller.

To get detailed information on the single devices, double click on the icons. The windows that open, depend on the type of the selected device.

To change the physical configuration of the devices, click the right mouse button on the desired device or select the corresponding item from the Configuration menu. Possible choices are (Click on the popup menu items to get information on them):

D.4.1 Controllers

Icon	Description	double click opens...	right click opens...
	Controller	Controller Information	 <ul style="list-style-type: none"> Change Settings Save Information Firmware Update Silence

Controller Settings

Type: Slot:

Features:

Processor:

RAM:

Firmware: Serial No.: Driver:

Cache: Delayed Write:

BIOS: BIOS Warning Level:

Supported BIOS Drives: Memory Test:

Controller Settings

Cache: Off On

Delayed Write: Off On

BIOS: Removed Enabled Disabled

BIOS Warning Level: All Messages Fatal Errors

Supported BIOS Drives: 2 7

Memory Test: No Test Standard Double Intensive

Change Settings

Here you can change the settings of the ICP Controller.

Cache	Enables or Disables the ICP Controller cache. For optimum performance the cache should be always On.
Delayed Write	Enables or disables the write cache function of the ICP Controller cache. For optimum performance the write cache should be always On.
BIOS	The BIOS of the ICP Controller is needed to boot the computer and the operating system from a Host Drive.
BIOS Warning Level	Enables or disables the display of non-critical boot messages of the ICP Controller during the system boot phase.
Supported BIOS Drives	Adjust the number of Host Drives which are supported by the ICP Controller's BIOS. This is the number (2 or 7) of Host Drives which are available under DOS without loading a special driver.
Memory Test	Configures the strategy and duration of the ICP Controller cache memory test during the boot phase. A more intensive test requires more time.

Save Information

After the selection of this option a file dialog is opened, which allows you to specify the path and name of the Save Information file. This file has a standard ASCII format and can be viewed or printed with a normal editor (e.g., notepad) or word processing system.

The Save Information file contains all relevant information on the ICP Controller (including firmware version, cache size, connector assignment, termination assignment), the connected devices (e.g., firmware version, SCSI parameters, selected transfer rate, number of grown defects, last status information), the Logical Drives , Array Drives and Host Drives. Thus, it represents a very easy and effective way to create a detailed documentation of the ICP Controller(s) and disk array(s).

The Save Information file can also be very helpful for a remote diagnosis of a system. If a trained support person has a copy of this file (by fax, by email), she/he can very easily find the appropriate steps to bring the system back to full operation.

Update the ICP Controller Firmware

The firmware, the BIOS and ICPCON of the ICP Controller are stored in a Flash-RAM which is part of the ICP Controller hardware. In contrast to EPROMs, Flash-RAMs can be re-programmed many times and without the complicated UV-light erasing procedure. Thus, both software modules can be easily updated without having to remove the controller from its PCI slot. Firmware and BIOS are part of the GDT_RPFW file.

To get the latest firmware for your ICP Controller, you can visit our website at

<http://www.icp-vortex.com>

or our ftp server at:

<ftp://ftp.icp-vortex.com/download>

The file has an extension (e.g., GDT_RPFW.009) which indicates the version stepping. We recommend that you also download the packed files which contain the latest programs/drivers for the operating system used on your system. Observe the following order when carrying out the updating procedure:

- Get the latest GDT_RPFW file for the ICP Controller. The file does NOT need to be expanded !
- Format a 3.5" HD disk (1.44MB) and copy the GDT_RPFW file on this disk.
- After selecting the Firmware Update option a file dialog opens where you should specify the path (A:). A list of firmware versions (normally only one) is displayed. After selecting the new firmware and final confirmation, the new firmware is programmed into the Flash-RAM of the ICP Controller.

The new firmware becomes active after the next cold boot.

Note: All user specific settings concerning the ICP Controller and the disk arrays are not affected by the firmware update.

Turn off the Audible Alarm

This option allows you to turn off the audible alarm of the ICP Controller manually. After a significant event (a drive failure or an overheat of the ICP Controller) the audible alarm of the ICP Controller is turned on. If the reason of the event is removed, the audible alarm turns off automatically.

If a member of an Array Drive has failed you should replace the failed drive as soon as possible. Read more about the fail state.

If the state of an Array Drive changes into error (more than one drive has failed) please contact our hotline for further assistance.

If no drive failed, you may also check the Controller Events to find out if the CPU of the ICP Controller is overheated.

If you hear an audible alarm, but the silence function is not offered and there are no critical controller events, check, if the alarm is coming from a different source, like the CPU cooler or the mainboard.

D.4.2 I/O Processors

These are the I/O processors of the ICP Controller. Each processor controls one I/O channel.

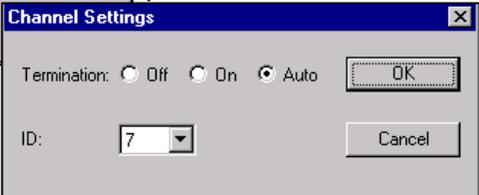
Icon	Description	double click opens...	right click opens...
	SCSI processor SE (single ended)	Channel Information	
	SCSI processor LVD/SE (low voltage differential / single ended)		
	FC processor FC AL (fibre channel arbitrated loop)		



Channel Info dialog box showing: Type: SCSI, Processor: 53C875, narrow/wide term.: 0n/0n, i68p/150p/e68p/e50p: 1/0/0/-, Termination: Auto, ID: 7.



Rescan ID(s) dialog box showing: Rescan all ID(s), Choose ID: 0.



Channel Settings dialog box showing: Termination: Off On Auto (selected), ID: 7.



Hot Plug ID dialog box showing: Choose ID: 0.

Channel Settings

Termination

The termination for this channel of the ICP Controller can be set to three different states:

- ON: The termination of the lower (low byte) and upper data lines (high byte) is enabled or disabled depending on the occupied SCSI connectors of this channel.
- OFF: No lines are terminated
- AUTO: All 8/16 data lines are terminated.

ID

Changes the SCSI ID of the SIOP (0-7).

Warning: The SCSI bus termination of The ICP Controller must match the existing SCSI cabling and cable termination. If the ICP Controller represents one end of the SCSI bus its termination must be ON or AUTO. If you change the termination to OFF, it is very likely that the SCSI bus is longer stable and the connected devices fail. Due to the different technology, these settings cannot be changed with an FC I/O processor (**fibre channel**).

Rescan ID(s)

This function allows you to rescan one or all IDs of the selected SCSI channel. It displays an overview of all powered SCSI devices which are currently connected with the SCSI bus.

Warning: If the hard disk contains valid Host Drive or Array Drive information, they will be deleted when creating a new Logical Drives. All data will be lost.

Hot Plug: Add Disk

With this Hot Plug function you can add a new hard disk while the system is running. Before starting the Hot Plug procedure, you should prepare the new hard disk and set the jumpers according to the free SCSI IDs and SCSI bus termination. Thus, pay attention to the correct termination of the SCSI bus and set the ID of the hard disk to a free address, if that is not done automatically by your storage subsystem. If you do not know, which IDs are in use on the selected SCSI channel, you can find this out in the physical configuration window. Wrong SCSI bus termination and/or SCSI ID conflicts will cause a failure of the complete SCSI channel. During the Hot Plug procedure firstly enter the ID of the new hard disk (all free IDs are offered). Secondly, you have to confirm a security request. Thereafter the SCSI channel is halted. While the channel is halted, you have 45 seconds to plug in the new hard disk. (Note: Since the operating system will timeout, the channel cannot be halted longer). If you do not complete the Hot Plug procedure manually, the channel will automatically start again after 45 seconds. After the successful completion of the Hot Plug, the new Physical Drive appears in the physical configuration window. Now, you can use it to build a new Host Drive, replace a failed drive or for the expansion of an existing Array Drive. You should only use the Hot Plug function with servers and/or subsystems which are designed and certified for hot plugging of hard disks.

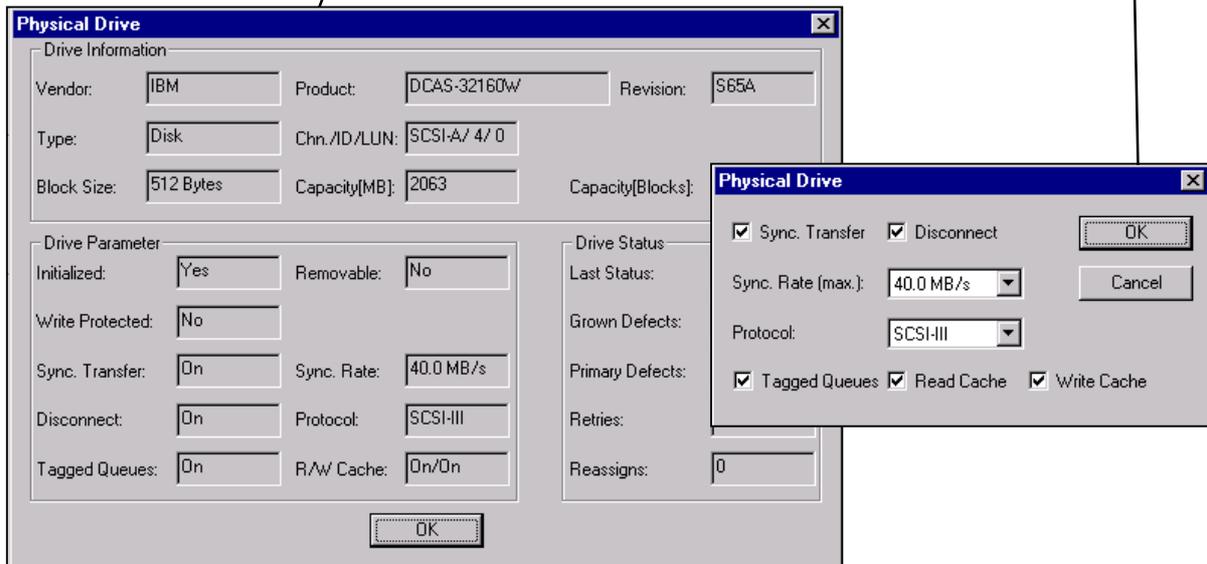
Warning: If the new plugged in hard disk contains valid Host Drive or Array Drive information, they will be deleted when creating a new Logical Drives. All data will be lost.

D.4.3 Direct Access Devices

Direct access devices can be initialized and subsequently used for Logical Drives. The ICP Controller caches the I/Os of these devices. Logical Drives are the components of Array Drives.

Note: If you want change the media of a removable disk during operation, the media MUST NOT BE INITIALIZED with ICP RAID Console (ICPCON) or the ICP RAID Navigator. Furthermore these devices have to be reserved for the raw service. Thus, the removable disk is handled like a non direct access device.

Icon	Description	double click opens...	right click opens...
	Physical Disk	Physical Drive Information	 <ul style="list-style-type: none"> SCSI Parameter/Initialize Format Disk Check Surface Deinitialize Disk Progress Information Lock Disk Unlock Disk Hot Plug: Remove Disk
	Removable Disk		



The SCSI Parameters

The SCSI parameters configure a Physical Drive (especially its performance). By changing the SCSI parameters, you can

- set the synchronous transfer rate
- change the settings of the hard disk cache
- enable or disable the disconnect feature
- enable or disable the tagged queues

Warning: By changing these parameters on a new hard disk or a hard disk, which has been connected with a non ICP Controller, this hard disk will be initialized and all data on this hard disk will be lost.

You can access the change SCSI Parameters / Initialize menu by clicking the right mouse button on a Physical Drive in the Physical Drives Windows (View > Physical Configuration)

Synchronous Transfer

Different to the older asynchronous transfer mode, the synchronous transfer offers higher transfer rates on the SCSI bus. The maximum synchronous data transfer rate depends on the width of the SCSI bus (8 bit narrow SCSI, 16 bit wide SCSI) and the frequency of the signals on the SCSI bus (10 MHz for Fast SCSI, 20 MHz for Ultra SCSI and 40 MHz for Ultra2 SCSI and Ultra 160 SCSI). The maximum data transfer rates are:

	Narrow SCSI (8Bit)	Wide SCSI (16Bit)	
Fast SCSI	10 MB/s	20 MB/s	
Ultra SCSI	20 MB/s	40 MB/s	
Ultra2 SCSI	-	80 MB/s	(LVD SCSI (l ow v oltage d ifferential) only)
Ultra 160	-	160 MB/s	(LVD SCSI (l ow v oltage d ifferential) only)

Note: The maximum synchronous data transfer rate between a SCSI device and the ICP Controller can be limited. This limitation may become necessary if a particular SCSI cabling does not allow the maximum rate the controller and the drive could achieve. The adjustment of the synchronous transfer rate can be done in the SCSI parameters / Initialize menu. Notes: In order to select a transfer rate above 10 MB/s the protocol has to be set to SCSI-III. Higher frequencies (single ended SCSI) require better cabling, shorter cables and a professional termination of the SCSI bus.

Disconnect

This SCSI feature enables a SCSI device to disconnect from the SCSI bus. By releasing the SCSI bus while the device doesn't need it, the bus becomes free for other devices to transfer or receive data. An example for a disconnect situation is a hard disk that needs time to read data from its media after receiving a read command. It may then disconnect from the SCSI bus so that other devices can transfer data. Later on it reconnects when it is able to deliver the data instantly. You can change the disconnect setting for a hard disk in the SCSI parameters / Initialize menu.

Tagged Queues

This feature enables modern hard disk drives to receive several commands at one time and then process them in an optimized order. Few, mainly older drives, do not or not correctly support this feature. You can change the tagged queues setting for a hard disk in the SCSI parameters / Initialize menu.

SCSI read cache / SCSI write cache

These settings enable or disable the write and read caches of the hard disks. Turning these caches off degrades the performance of the hard disk. You can change the cache settings of a drive in the SCSI parameters / Initialize menu.

Format a Physical Disk

Caution: This function deletes ALL DATA on the selected Physical Drive !

The hard disk which you want to low level format, may not be member of a Logical Drive if you want to start a low level format or a check surface. This function sends a format unit command to the Physical Drive. Everything else is done by the drive itself. The ICP Controller stands by to receive a good status back from the drive when it has successfully finished the low level format, thus no progress information can be shown. What the drive exactly does during the low level format and how long it takes depends on the manufacturer's format unit command implementation in the firmware of the drive, the speed and capacity of the drive. A low level format can take from seconds to hours or even days. Before you confirm to start the low level format, you are asked if you want to delete the grown defect list of the drive. Deleting this list, does not mean that the grown defects (media defects) vanish. It is very likely that the low level format will trace them again and map them out. Usually, there is no need to low level format a drive. This should be only done with a defective drive. If you want to check the surface of a hard disk, it is better to use the check surface function.

Warning: It is strongly recommended not to interrupt a low level format of a drive. This may cause an inoperable drive.

Deinitialize a Physical Disk

Deinitializing a Physical Drive removes the ICP initialization sectors. The Physical Drive may not be member of a Logical Drive if you want to deinitialize it.

Lock / Unlock a Removable Disk

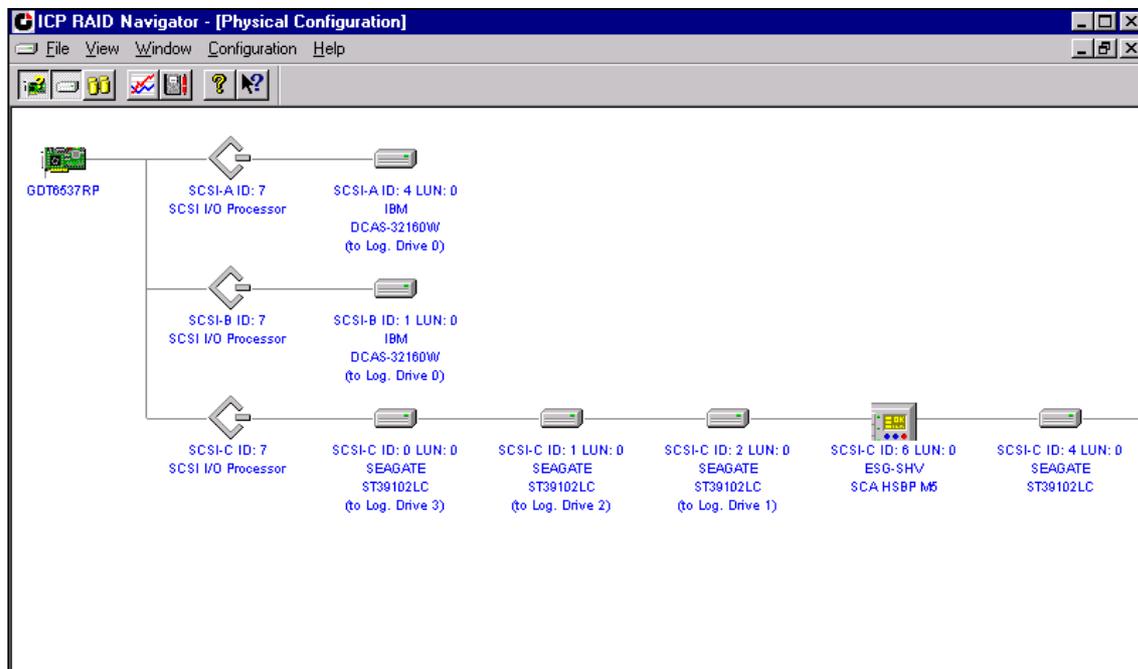
With this function you can lock or unlock a removable drive containing a removable medium. If the Physical Drive is locked, the medium cannot be ejected with the eject button while the system is running.

D.4.4 Non direct access devices (raw devices)

Non direct access devices cannot become components of Logical Drives, Array Drives or Host Drives. These devices are either controlled by a software driver (e.g. an ASPI module), the operating system or an application. Non direct access devices cannot be initialized or changed in their SCSI parameters with this program.

Icon	Description	double click opens...	
	CDROM / CDR	Physical Drive Information	
	Tape drive		
	SAF-TE-Processor		
	Scanner		
	SCSI Printer		

Example for a Physical Configuration and the Controller Settings:



D.5 Logical Configuration Window



This window shows the logical configuration of Host Drives, Array Drives and Logical Drives controlled by the selected ICP Controller. The complete configuration is shown as a tree starting from the left with the Host Drives, followed by the Array Drives (if RAID is configured), the Logical Drives and the Physical Drives. To get detailed information on the single devices, double click on the icons. The windows that will open then, depend on the type of the device.

To change the logical configuration of the drives, click the right mouse button on the desired drive or select the suitable option from the configuration menu.

Host Drives

These are the drives "visible" to the operating system.

Icon	Description	double click opens...	right click opens...
	Host Drives	Host Drive Information	<div style="border: 1px solid gray; padding: 5px; width: fit-content;"> <p><u>D</u>rive Name</p> <hr/> <p><u>C</u>reate New</p> <p><u>R</u>emove</p> </div>
	Create a new Host Drive	Create a new Host Drive	



Normal Host Drive

This Host Drive belongs to the selected ICP Controller on a non clustering I/O channel.



Local mounted Cluster Drive

A Cluster Drive is a Host Drive on a clustering I/O channel.

This Host Drive is mounted on the selected ICP Controller.

Clustering means that two or more servers share resources. In case of one server failing these resources move over to another valid server. The hard disks used for clustering are physically not connected with just one ICP Controller, but with several ICP Controllers in different servers (shared I/O channels).

Only one server can own a Host Drive at the same time, so a Host Drive configured for clustering may either be controlled by the selected ICP Controller (i.e., local mounted) or by a different ICP Controller (i.e., remote mounted).

You may only change the configuration of a local mounted Host Drive.



Remote mounted Cluster Drive

A Cluster Drive is a Host Drive on a clustering I/O channel.

This Host Drive is mounted on an ICP Controller in a different server.

Clustering means that two or more servers share resources. In case of one server failing these resources move over to another valid server. The hard disks used for clustering are physically not connected with just one ICP Controller, but with several ICP Controllers in different servers (shared I/O channels).

Only one server can own a Host Drive at the same time, so a Host Drive configured for clustering may either be controlled by the selected ICP Controller (i.e., local mounted) or by a different ICP Controller (i.e., remote mounted).

You may only change the configuration of a local mounted Host Drive.



Private Host Drive belonging to another ICP Controller

This is a Host Drive on a clustering I/O channel.

This Host Drive is configured as a Private Host Drive and does not belong to the selected ICP Controller.

Clustering means that two or more servers share resources. In case of one server failing these resources move over to another valid server. The hard disks used for clustering are physically not connected with just one ICP Controller, but with several ICP Controllers in different servers (shared I/O channels).

Usually, when booting a cluster, the ICP Controller starting up first will gain control over all Host Drives on the clustering I/O channels. If you want to be sure that a Host Drive is controlled by a specific server, you may configure this Host Drive as a Private Host Drive. A Private Host Drive cannot be used for clustering anymore.

Note: You cannot change the configuration of a Private Host Drive which is not property of the selected ICP Controller. You have to use ICPCON to change its configuration.

A Private Host Drive belonging to the selected ICP Controller will be displayed like a normal Host Drive.

Array Drives

Array Drives or RAID drives consist of Logical Drives. They can be fault tolerant, depending on the RAID level. The RAID level is displayed under the Array Drive icon.

Click on the icons to get more information on the different states of an array.

Icon	Description	double click opens...	right click opens...
	RAID 4/5/10	Array Drive Information	<ul style="list-style-type: none"> Drive Name Parity Verify Parity Recalculate Expand Array Progress Information Add Hot Fix Remove Hot Fix Hot Fix Pool Access Add RAID1 Component Remove RAID1 Component
	Idle		
	Build		
	Ready		
	Fail		
	Error		

  	RAID 1		
	Build		
	Ready		
	Fail		
	RAID 0		

Logical Drives

Logical Drives consist of one or more Physical Drives.

Icon	Description	double click opens...	right click opens...
	Single Disk	Logical Drive Information	
	Hot Fix Drive		
	Failed / missing Hot Fix or Single Disk		
	Chaining Drive		

Physical Drives

These are the hard disks. You cannot change any settings here. If you want to change the settings, you have to do this in the physical configuration window.

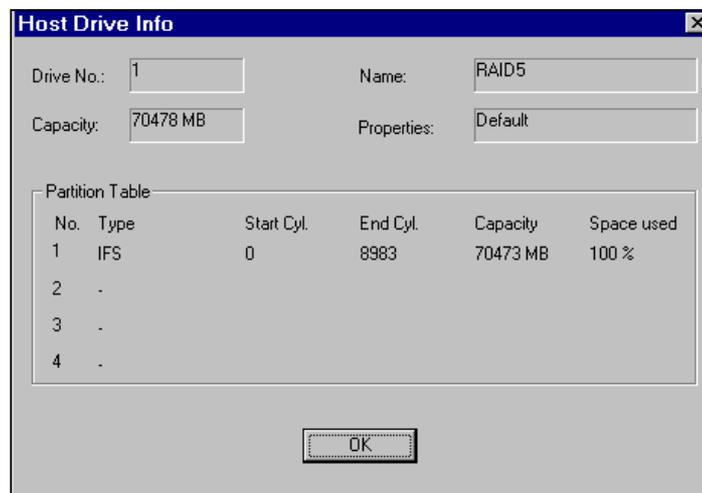
	Physical Disk	You can't change the disk settings here. Use the Physical Configuration window instead.
---	---------------	---

D.5.1 The Host Drive Information Window

Double click on the Host Drive icon.

This window contains information on a Host Drive like the Host Drive's capacity and a possible partition table.

Drive No	The Host Drive number of the Host Drive. The Host Drives are reported to the system one after the other, beginning with the lowest Drive Number. If the corresponding ICP Controller is the first controller in the system, the system will boot from the Host Drive with the lowest number.
Name	The name of the Host Drive. A Host Drive is automatically given a name during the configuration. It can be renamed by selecting the menu Drive Name of the menu list you get after clicking with the right mouse button on the Host Drive icon.
Capacity	The capacity reported to the operating system. The size of a Host Drive can be smaller than the size of the Logical Drive or the Array Drive the Host Drive is a part of.
Partition Table	The partition table of a Host Drive can contain up to four partition entries. Information on the type of the partition, it's location on the Host Drive, it's size and the percentage of space used by this partition are shown.



D.5.2 The Array Drive Information Window

Double click on the Array Drive Icon.

This window shows information on an Array Drive (i.e., a disk array).

Drive No	The logical number of the Array Drive.
Drive Name	The name of the Array Drive.

Type	This is the RAID level of the Array Drive. It can be RAID 0 (data striping, no redundancy), RAID1 (mirroring), RAID 4 (striping with parity drive), RAID 5 (striping with striped parity) or RAID 10 (combination of RAID 1 and RAID 0).
Status	The status of an array can be ready (operational), fail (one drive missing, still operational but not redundant), error (more than one drive failed, not operational), rebuild (one drive is being integrated), build (initial preparation of the array), idle (new defined array, before build process has started), expand (one or several drives are added online to the Array Drive).
Capacity	This is the capacity available for the corresponding Host Drives.
Attributes	The attribute of an Array Drive is usually read/write ([RW]). If a component of a disk array is missing while startup and the operator decides not to activate fail mode, the array is set to the read only attribute ([RO]). When the missing drive is available again, the Array Drive shows again the ready status (i.e., attribute [RW]).
Striping Size (RAID 0, 5, 10)	Data written on RAID 0, 4, 5 or RAID 10 drives is distributed over all drives (striping). This is the size of the blocks. The striping size can only be changed when the drive is created (standard setting is 32KB).
Invalid/Missing (RAID 1)	These two values show the number of invalid/missing drives. While the build is in progress (initial copy of data from the Master to the Slave), invalid drives exist (the drives are updated).
Pool Hot Fix	Indicates, if the Pool Hot Fix access is enabled or disabled. 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 Array Drives as a Hot Fix drive (assuming it has an appropriate capacity). 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. Hot Fix Drives can also be configured as Private Hot Fix Drives. A Private Hot Fix Drive can only be used by the Array Drive it was configured for. The Pool Hot Fix access can be changed by clicking with the right mouse button on the Array Drive icon.

RAID4/5/10 Drive Info

Drive No.:	<input type="text" value="1"/>	Drive Name:	<input type="text" value="RAID5"/>
Type:	<input type="text" value="RAID-5"/>	Status:	<input type="text" value="ready"/>
Capacity:	<input type="text" value="70478 MB"/>	Attributes:	<input type="text" value="RW"/>
Striping Size:	<input type="text" value="32 KB"/>	Pool Hot Fix:	<input type="text" value="Disabled"/>

D.5.3 The Logical Drive Information Window

Double click on the Logical Drive icon.

This window shows information on a Logical Drive. A Logical Drive can be either a single disk, or a chaining group of disks (concatenation), or a stripe set of several disks.

Drive No	The number of the Logical Drive.
Drive Name	The name of the Logical Drive.
Type	The type of the Logical Drive. A Logical Drive can be either a single disk, or a chaining group of disks (concatenation), or a stripe set of several disks.
Status	This is the state of the Logical Drive. It can be ready, missing (not available after reset) or failed (i.e., no longer available for the controller).
Capacity	The capacity which is available for Array Drives or Host Drives.
Attributes	The attribute of a Logical Drive is usually read/write ([RW]). If a component of a disk array is missing while startup and the operator decides not to activate fail mode, the array is set to the read only attribute ([RO]). When the missing drive is available again, the Array Drive shows again the ready status (i.e., attribute [RW]).
Last Status	The last status information of a Logical Drive. This is different to the last status information of Physical Drives.
RAID 1 Component	If the Logical Drive is a member of a RAID 1 array, it can be either master or slave. If a RAID 1 component is added to a Logical Drive, the data is read from the master and written to the added Logical Drive (the slave). One master may have several slaves.
RAID 4/5 Component	As a member of a RAID 4 or RAID 5 array a Logical Drive can be the master or a component. Different to RAID 1 arrays, this is only important for the internal structure of the array.

The screenshot shows a window titled "Log. Drive Info" with the following fields:

- Drive No.: 1
- Drive Name: DISK_A1
- Type: Disk
- Status: ok
- Capacity: 35239 MB
- Attributes: RW
- Last Status: 0
- RAID1 Component:
 - Type: ---
 - Status: ---
- RAID4/5 Component:
 - Type: Master
 - Status: ok

An "OK" button is located at the bottom of the window.

D.5.4 Change the name of a Drive

Click the right mouse button on the drive icon.

This menu option opens a dialog where you can enter a new name for the selected drive (up to 7 characters).

D.5.5 Remove a Host Drive

Click the right mouse button on the Host Drive icon. By selecting this menu item you can delete a Host Drive.

CAUTION !

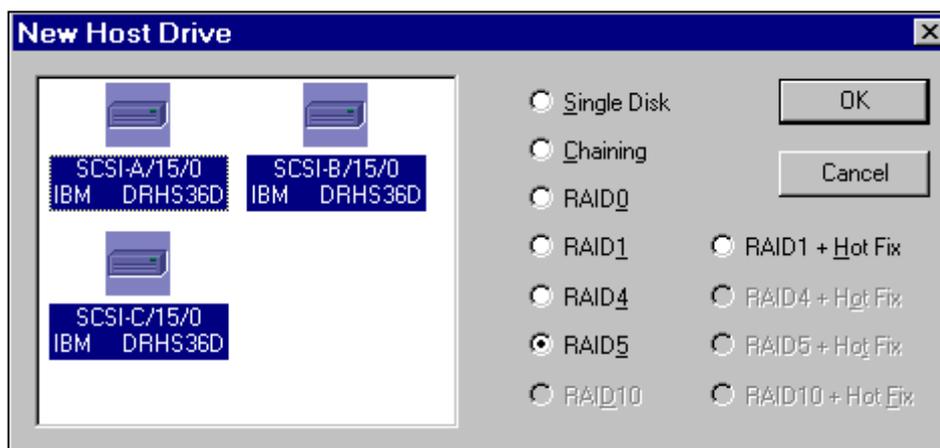
By deleting a Host Drive, all other Host Drives which are part of the same Array Drive or Logical Drive, the Array Drive or Logical Drive itself and all its components are deleted ! ALL DATA WILL BE LOST.

To prevent the operating system from crashing and to hinder the unintentional deletion of drives all partitions on the Host Drive (and it's related Host Drives if it is split into several Host Drives) must be deleted before the Host Drive can be removed.

D.5.6 Create a new Host Drive

Double click on the new Host Drive icon. Only the Host Drives are "visible" for the operating system.

The structure of the Host Drives is not known to the operating system i.e., the operating system does not recognize that a given Host Drive consists of a number of hard disks forming a disk array. To the operating system this Host Drive simply appears as one single hard disk with the capacity of the disk array. This complete transparency represents the easiest way to operate disk arrays under any operating system, neither operating system nor the computer need to be involved in the administration of these complex disk array configurations. To create a new Host Drive, you need one or more Physical Drives which are not yet part of another Host Drive. After double clicking on the New Host Drive icon a new window opens.



On the left side you see a box with the available Physical Drives, under the drives you can see their physical coordinates (channel/ID/LUN), the manufacturer and the vendor-unique name. On the right side you see a list of all possible Host Drive types which can be configured. When there is no Physical Drive selected in the left box, all possibilities are disabled

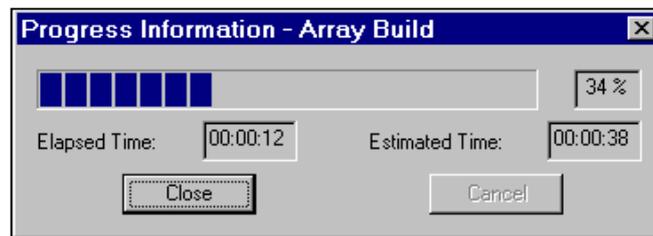
(gray). While you are selecting Physical Drives for the new Host Drive, more and more possibilities become selectable.

You can select a single Physical Drive by clicking on it. If you want to select more than one Physical Drive, simply draw a frame around the Physical Drives, or press the <ctrl> key and then click on all the Physical Drives you want to combine to a new array.

When you have finished the selection of the Physical Drives, choose the type of Host Drive you want to create and click on OK.

If the selected Host Drive uses data striping (RAID 0, 4, 5 or 10) you can change the default striping size. If you have selected a configuration with Hot Fix Drives, you can choose between a Private Hot Fix Drive or a Pool Hot Fix Drive. After the Host Drive was created, you can partition and format the Host Drive with the corresponding operating system utility.

If an array build started, you can monitor the progress of the array build by clicking the right mouse button on the Host Drive and then selecting progress information.



D.5.7 Parity Verify

Click the right mouse button on the Array Drive icon.

RAID 4 and RAID 5 drives contain parity information, which is used in case of a drive failure. The parity information is calculated from the user data on the disk array. On RAID 4 disk arrays the parity data is stored on a single disk (parity disk), on RAID 5 disk arrays the parity data is being distributed over all drives (parity striping). This option verifies online the parity information of the selected RAID 4 or RAID 5 Array Drive. If this option is selected for several Array Drives, the processes are put into a queue and performed one after the other.



If a parity error is detected, you should try to find the reason for this data corruption. A good indication for data corruption can be retries on the SCSI bus. If the retry-counter shows high numbers, this might be the problem. Possible reasons for parity error are bad cabling or termination or a hardware error like a defective drive or a drive which is overheated. **After** removing the reason of the data corruption you can carry out parity recalculate to ensure that the parity information of this disk array becomes again valid.

D.5.8 Parity Recalculate

Click the right mouse button on the Array Drive icon.

A parity recalculate can be used to repair parity errors which have been previously detected with a parity verify. A parity recalculation initiates the same process as used for a build on

an Array Drive. The user data on the drives is read, parity information is calculated from this user data and the parity information is written anew.

While the parity recalculate is in progress the array is in the build state. During this time the Array Drive is not redundant. You can view the progress of the build process by when you click the right mouse button on the Array Drive icon and select progress information.

After the parity recalculate is completed the Array Drive's state changes again into ready with the addition 'patch'. This notification has no relevance for the operation of the Array Drive but is a reminder that the parity information of this Array Drive has been recalculated once.

D.5.9 Progress Information

Click the right mouse button on the Array Drive or Physical Drive icons.

The progress information window shows the progress (elapsed time, estimated time, percentage of completion) of a parity verify, a disk array build or a surface check. During an Array Drive rebuild, the information (user data and/or parity) of a specific Logical Drive is calculated from the user data and parity information of the other Logical Drives and written to this new Logical Drive. During a build of an Array Drive parity information is calculated and written to a specific Logical Drive (RAID 1/4) or striped over the Array Drive (RAID 5). The estimated time indicates the time needed by the controller to finish this process if the user load on the controller does not change. The elapsed time is reset when opening the window.



D.5.10 Expansion of an Array

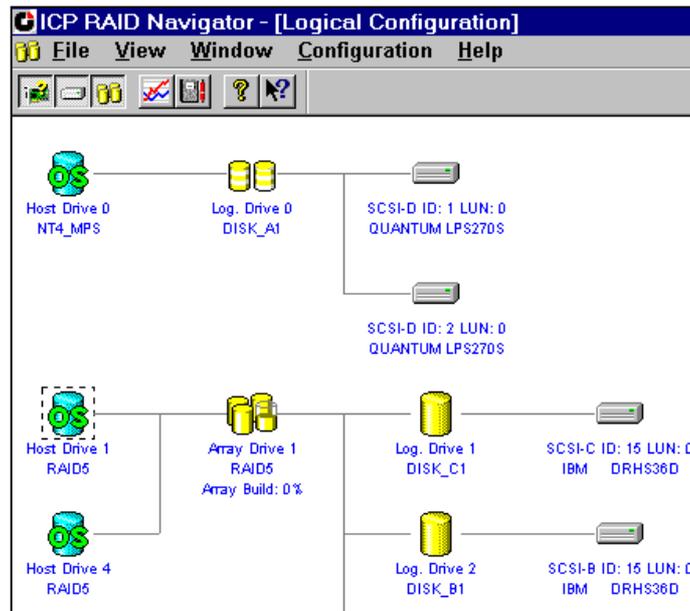
Click the right mouse button on the Array Drive icon.

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) and/or

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.

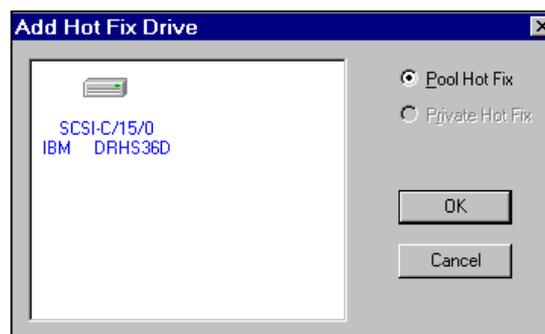


If you want to add additional drives to the Array Drive, select them from the box on the left side of the window. It is possible to add more than one drive at the same time. If no Physical Drives are offered, you have to use the Hot Plug: Add Disk function first, to add new drives. You can open the progress information window to monitor the progress of the expansion. If you have previously limited the capacity of the Array Drive, you are asked, if you want to convert this free space to a new Host Drive.

IMPORTANT: Before starting an expansion it is absolutely necessary to verify that you have a valid backup of the complete data on the system. The expansion process includes a new, unknown hard disk. If there are problems with this disk or with the cabling of this disk there is always the risk of data loss. Some expansions will take quite a long time, so it is advisable that there is not too much traffic (i.e., user load) on the system, otherwise an expansion of a large Array Drive can easily take days.

D.5.11 Add a Hot Fix Drive

Click the right mouse button on the Array Drive icon. Use this option to add a Hot Fix Drive to an Array Drive (RAID 1/4/5/10). You can choose if you want to add a Private Hot Fix Drive or a Pool Hot Fix Drive. Private Hot Fix Drives are assigned to a specific Array Drive and are activated if a member of this Array Drive fails. Pool Hot Fix Drives can be used by any Array Drive with enabled Pool Hot Fix Access. RAID 1 Array Drives allow only Pool Hot Fix Drives. After choosing the type of Hot Fix Drive you can select a Physical Drive from the box which shows all suitable drives.



The capacity of the Hot Fix Drive has to be larger or equal than the capacity of the components of the Array Drive (Example: The Logical Drives of the Array Drive have 4.2GB capacity, thus the capacity of the Hot Fix Drive has to be 4.2GB or larger).

The spindle motor of the Hot Fix Drive is normally (i.e., when the Hot Fix Drive is not needed) stopped. Thus, it may take a few seconds until you get a response, if you click on that drive.

Some general notes on Hot Fix Drives.

One of the reasons that have led you to choose RAID disk arrays definitely lies with the redundancy, that is, the data security you still preserve even in the event of disk failure, thus resting assured against loss of data and time. Hot Fix Drives are possible with all RAID 1, 4, 5 and 10 disk arrays. In order to assist the following considerations, we define the term time without redundancy, TWR. Set apart the time needed to set up the disk array (state build), the time without redundancy should be kept as short as possible. Let us assume that one of the hard disks of the RAID 5 disk array we set up with example 1 fails. The disk array is without redundancy. TWR starts to run. Any superfluous prolongation of the TWR (because you have to get a replacement drive, or because you did not realize the failure immediately since you didn't hear the ICP Controller's alarm signal, or because nobody checked the file server) increases the risk of data loss which will occur if a second drive should fail. Therefore, new redundancy should be created as soon as possible and in an entirely automated manner. Integrating a Hot Fix Drive as an immediately available and auto-replacing drive is the only way to keep the TWR as short as possible. Only a Hot Fix drive can ensure optimal disk array security and constant data availability. Of course a Hot Fix Drive is not compulsory. If you control the disk array at regular intervals and immediately replace a defective drive (by shutting down the system or hot-plugging), you can do without a Hot Fix Drive.

D.5.12 Remove a Hot Fix Drive

Click the right mouse button on the Array Drive icon.

This option opens a box showing available Hot Fix Drives. Here you can select the Hot Fix Drive you want to remove from the Array Drive and then confirm your selection. You can remove any Pool Hot Fix Drives or the Private Hot Fix Drive of the selected Array Drive.

D.5.13 Hot Fix Pool Access

Click the right mouse button on the Array Drive icon.

Here you can enable or disable the access of an Array Drive to the pool hot fix of Hot Fix Drives. If the access is enabled this means that if a member of an Array Drive fails, a drive can be taken from the Hot Fix Pool and build automatically into the Array Drive. To be able to activate this feature, there must be suitable Logical Drives in the Hot Fix Pool. You can add drives to the Hot Fix Pool with the function Add Hot Fix Drives. By doing so, the Hot Fix Pool access for this specific Array Drive is activated automatically. For all other Arrays Drive you have to activate the access manually.

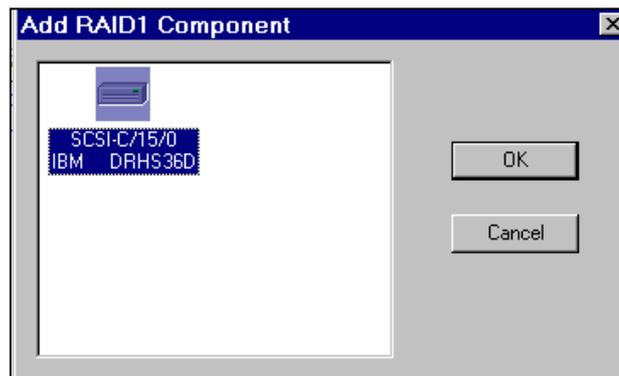
D.5.14 Add a RAID 1 Component (Mirror a Drive)

Click the right mouse button on the Logical Drive icon.

In certain "emergency" cases this is a very powerful and helpful option. This function allows you to add to a Logical Drive which is member of an Array Drive, another Logical Drive as a mirror drive (RAID 1).

Example: You have configured an Array Drive with 4 Logical Drives. One Logical Drive has failed and the Array Drive went into the fail state. Another failure would cause data loss. Unfortunately, you find another Logical Drive, which is shortly before failing (e.g., you hear a strange noise from it, or it's grown defect counter explodes). If you now initiate a hot plug it is very likely that this critical Logical Drive will also fail. This would result in a disaster.

To avoid that problem, you can mirror in a first step a new good Logical Drive to the critical one. When the copying is finished you remove the critical Logical Drive and then carry out a hot plug procedure.



To add a RAID 1 component you need to select in a first step a Logical Drive, the master. All suitable Logical Drives (with a capacity equal or larger than the capacity of the master) are shown in a box. Now choose the Logical Drive you want to add, the slave. After confirming, a new RAID 1 Array Drive is built, i.e., all data from the master are copied to the slave drive(s). When this process is finished all Logical Drives in the RAID 1 Array Drive are identical (there is no longer a difference between master and slave).

D.5.15 Remove a RAID 1 Component (Remove a Mirror Drive)

Click the right mouse button on the Logical Drive icon.

Here you can remove a drive from a RAID 1 Array Drive. If the Array Drive is in the build or in the fail state, you may only remove invalid RAID 1 components.

To remove a Logical Drive from a RAID 1 Array Drive you have to select the Logical Drive you want to remove. This is done by selecting a Logical Drive from a box where you see all members of the RAID 1 Array Drive. After confirmation the Logical Drive is removed. If the RAID 1 Array Drive consists of two drives (like usual) the RAID 1 Array Drive disappears and only one Logical Drive will be left. You can create a new RAID 1 Array Drive by selecting add RAID 1.

D.5.16 Replace a Logical Drive

Click the right mouse button on the Logical Drive icon.

If a Logical Drive of an Array Drive without a Hot Fix Drive should fail (or is very likely to fail, soon), you should replace the defective hard disk with a new one as soon as possible, because the Array Drive is without redundancy. The replacement Logical Drive has to have at least the same capacity as the failed one. The replacement is carried out with ICPCON or the ICP RAID Navigator. Before the Logical Drive can be removed, you have to select a new Logical Drive from the box with available Physical Drives which is shown after this option is selected. If no Physical Drive is offered, you have to use the Hot Plug: Replace Drive function to add a new drive. After the confirmation, the old Logical Drive is removed. Next, the data is rebuilt on the new Logical Drive. During this process the array is in the rebuild state and therefore not redundant.

D.5.17 The Different States of an Array Drive

The ready state



RAID 1



RAID 4/5

The Array Drive is fully operational when in the ready state. All redundant information is present, that is, a hard disk can fail without impairing the functionality of the Array Drive. This is the normal state of an Array Drive. The state ready/expand indicates, that the RAID level and/or capacity are currently migrated/expanded.

The idle state



RAID 4/5

This state is characterized by the fact that the redundant information of the Array Drive has never been entirely created. The Array Drive is in this state after its first configuration. If an error should occur while the array is in the build state, the array returns to the idle state (exception: if during build mode the dedicated drive of a RAID 4 Array Drive fails, the state changes to fail).

The build / rebuild state



RAID 1



RAID 4/5

After the Array Drive has been configured for the first time, and the build process is started it assumes the build state. While the Array Drive is in the build state, redundancy information is calculated and stored to the components of the Array Drive.

The disk array will assume the rebuild state after the automatic activation of a Hot Fix Drive or after a manual replacement (Hot Plug). The data and the redundant information are reconstructed and stored to the new drive.

In both states, the Array Drive is not redundant.

You can monitor the progress of the array build/rebuild by clicking the right mouse button on the Host Drive and then selecting progress information.

Note: User traffic on an Array Drive which is in the build/rebuild state, slows down the build/rebuild process.

The fail state



RAID 1



RAID 4/5

The Array Drive changes to the fail state whenever a Logical Drive fails. Redundancy information is still present, thus allowing the remaining hard disks to continue working. This state should be eliminated as soon as possible by replacing the defective hard disk. This can be done by using a Physical Drive, which is already connected with the controller, but not yet used for a Logical Drive, with the replace drive function, or by using the Hot Plug Replace Drive function. If a Hot Fix Drive has previously been assigned to an Array Drive, the ICP Controller will automatically replace the defective drive and start the reconstruction of the data and the redundant information. Therefore, under these circumstances the fail state is only temporary and will be eliminated by the controller itself.

Whenever an Array Drive enters a fail state, the ICP Controller's audible alarm is turned on. You can silence the audible alarm in the physical configuration window.

To analyze the reason for the drive failure, the last status from the Physical Drive information is very helpful. Additionally you should check for retries and/or reassigns.

A drive failure may also be the result of bad cabling, wrong termination or overheating.

The error state



RAID 4/5

If a second hard disk should fail while the Array Drive is in the fail or rebuild state, it is not possible to continue the working session without restrictions. The disk array is still available for I/Os, but data loss and error messages on the host level are possible.

Usually you have to remove the Array Drive and build a new one. In some situations (see below) there might be still a chance to reset the array. Please contact our support departments.

To find out why the drives failed, the last status from the Physical Drive information is very helpful. Additionally you should check for retries and / or reassigns.

A drive failure may also be the result of bad or cabling, wrong termination or overheating.

Some of these states may become the addendum patch (e.g. build/patch, ready/patch). This word indicates that the original Array Drive went through a significant procedure. I.e., the parity information was recalculated anew. Or, the Array Drive has been patched from the error state into the fail state. This may become extremely helpful in a situation where two Logical Drives of an Array Drive, fail at the same time, but only one of the two Logical Drives is really defective and the other was blocked out, since it was connected with the same I/O channel as the defective one. The Array Drive's state is error and normally all data would be lost. The ICP Controllers include some functions, which allow the patch of this Array Drive from the error state into the fail state. Before the actual patch, the defective drive has to be physically removed from the Array Drive. Such a patch-procedure is a real sheet-anchor and should only be used, after a detailed consultation with a trained support person (a printout of the Save Information file, is extremely helpful).

D.6 The Statistics Window



The statistics window can display the throughput of Physical, Logical and Host Drives. The vertical axis show the throughput, the horizontal axis the time. You can add drives by drag and drop them from the physical and logical configuration windows into the statistics window. If you want to change the layout of the lines, adjust the scales of the axis or remove drives from the statistics windows, you can do this using the chart menu.

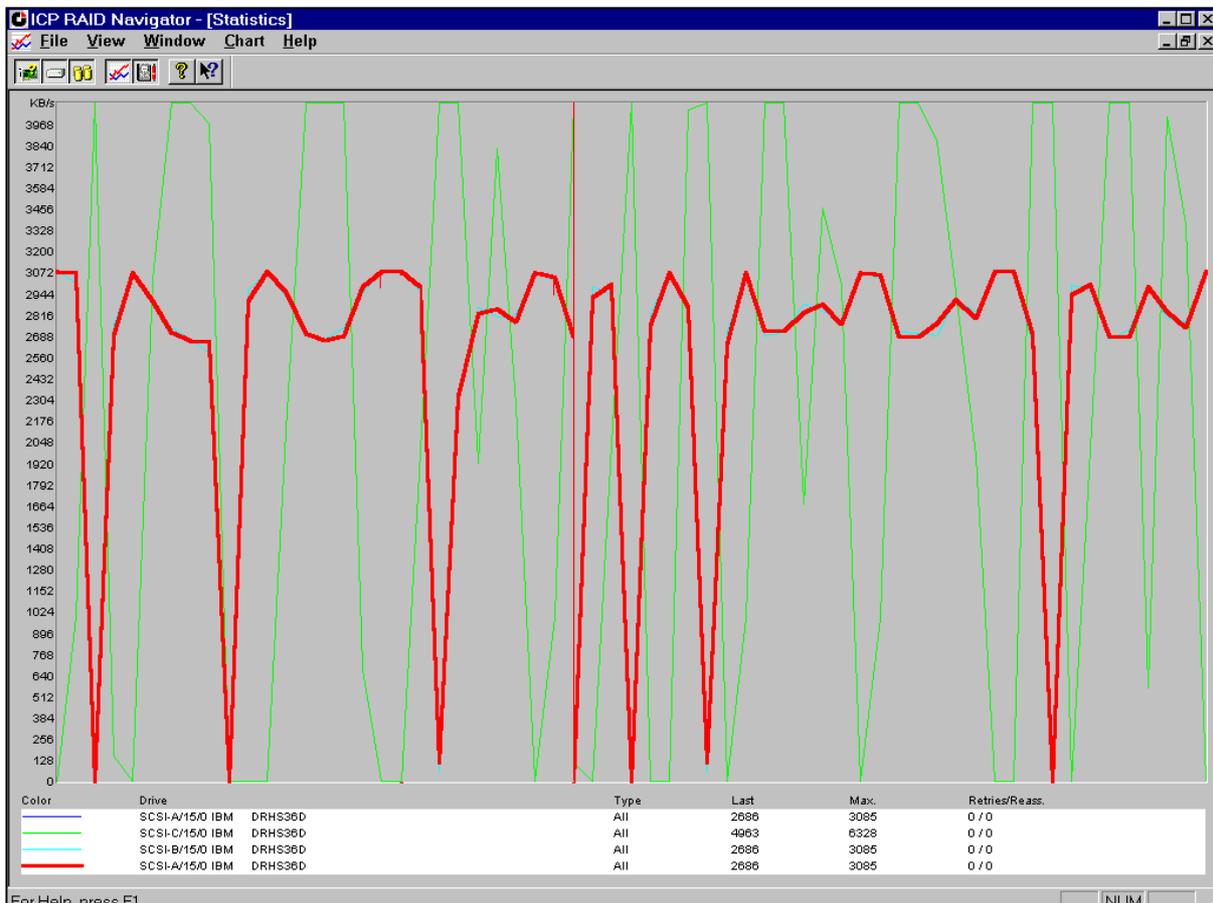
You can scale the statistics window using the cursor keys:



scale time down/up



scale throughput



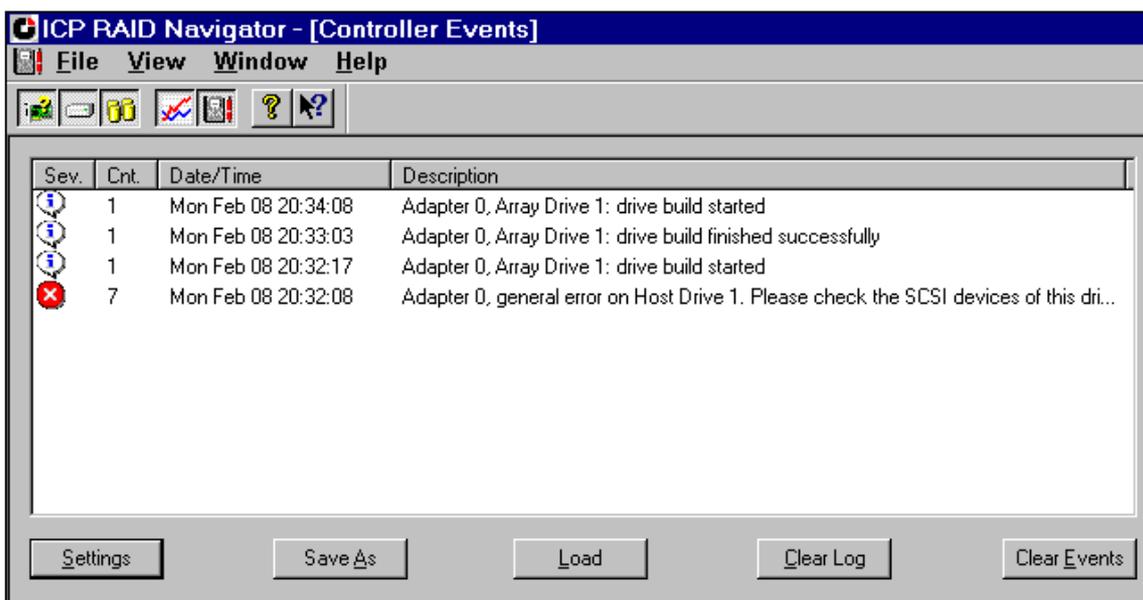
D.7 The Controller Events Window



This window shows the history of the events that occurred since the log was cleared the last time. The first column in this window contains icons representing the severity of the events:

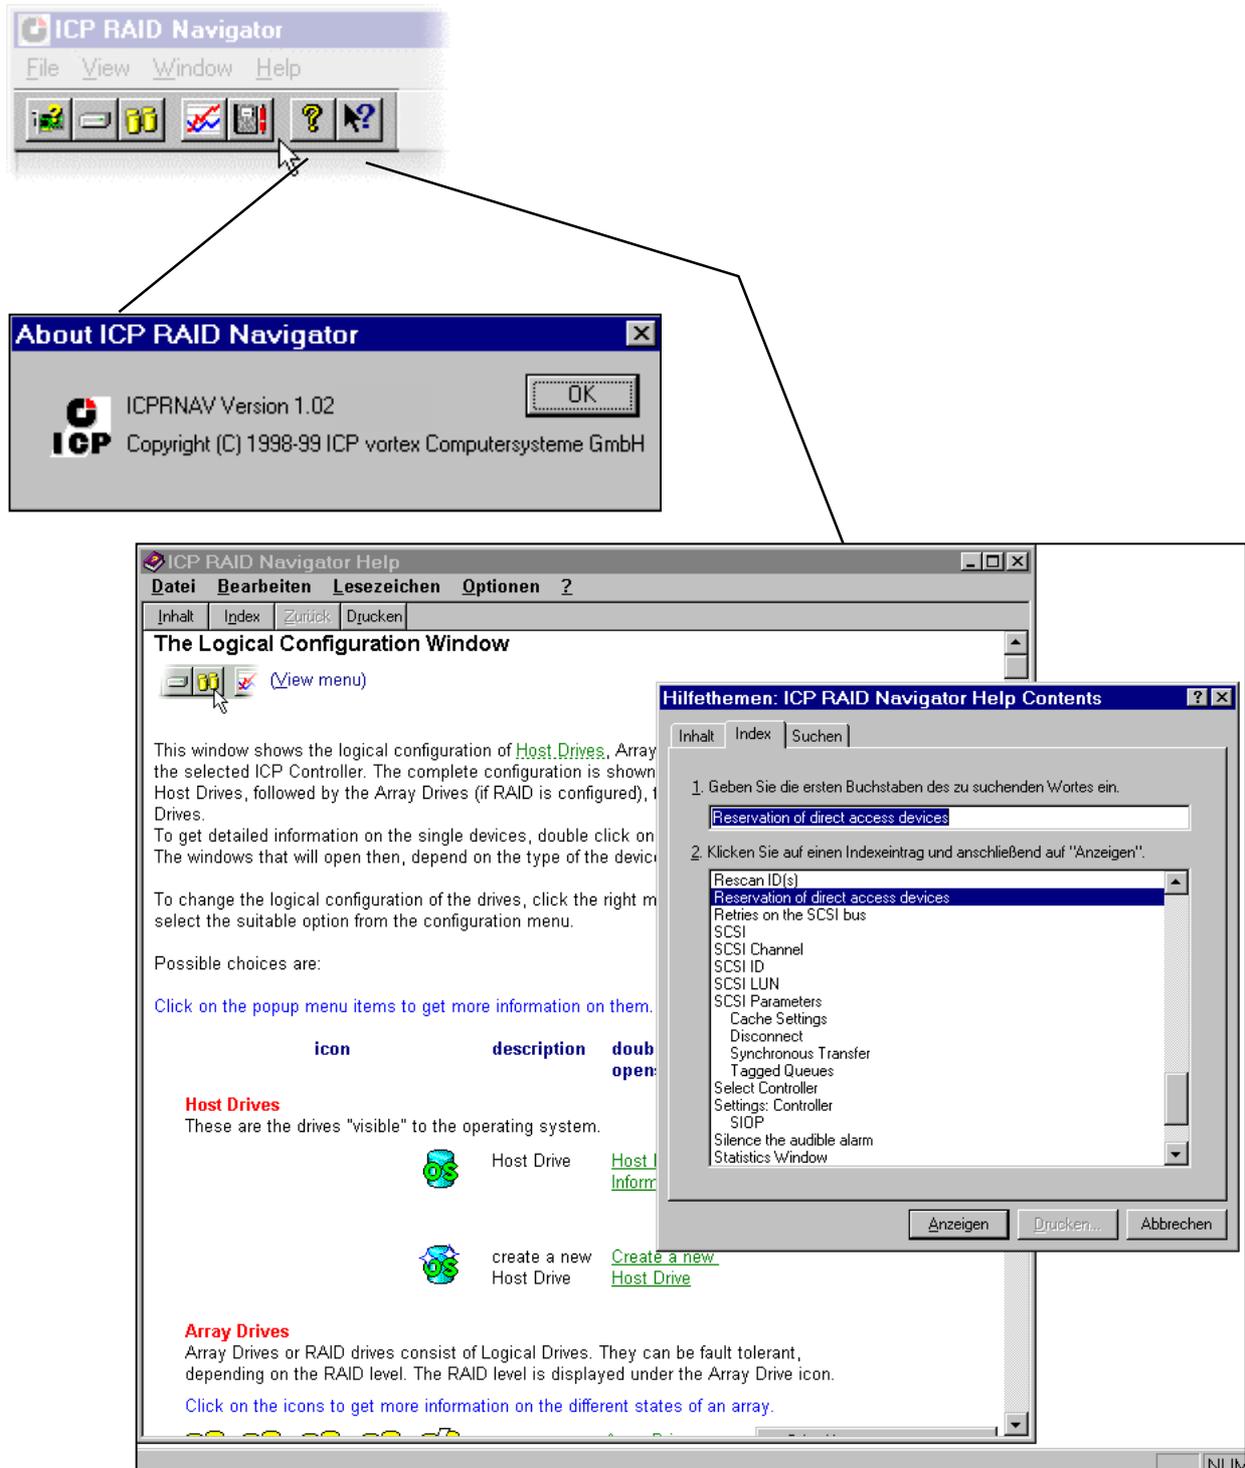
Icon	Description
	Information: This event is not critical. It only informs the system administrator about certain events, like the completion of an array build.
	Warning: This event is critical and may not be ignored. It occurs for example if an Array Drive's state changes from ready to fail.
	Error: An error occurred. This might be the failure of a drive.

The second column lists how often this event occurred. Some events may occur quite often, so this counter helps to keep an overview. The next column lists the date and the time when the event occurred. The last column describes the event. At the bottom of the window there is the Settings button. Here you can change the polling interval and enable the auto save function. Logs that are not saved on the hard disk are lost when the system is reset. The auto save function saves the current log in the time intervals selected there. The default name for this logfile is 'gdtevt.evt'. The Save As button allows you to write the contents of the event log to a specific file on the hard disk. The Load button loads an older event log and displays it. The last two buttons delete the log file and clear the event buffer on the ICP Controller.



D.8 ICP RAID Navigator Help

The ICP RAID Navigator includes an online help function. You can either choose the Help menu or the pointer with the question mark to obtain online help on a specific icon or function. There is also an index which allows you to search for certain keywords and/or topics.



D.9 ICP Service and ICP Mail

There are further powerful tools which are part of the ICP RAID Navigator delivery:

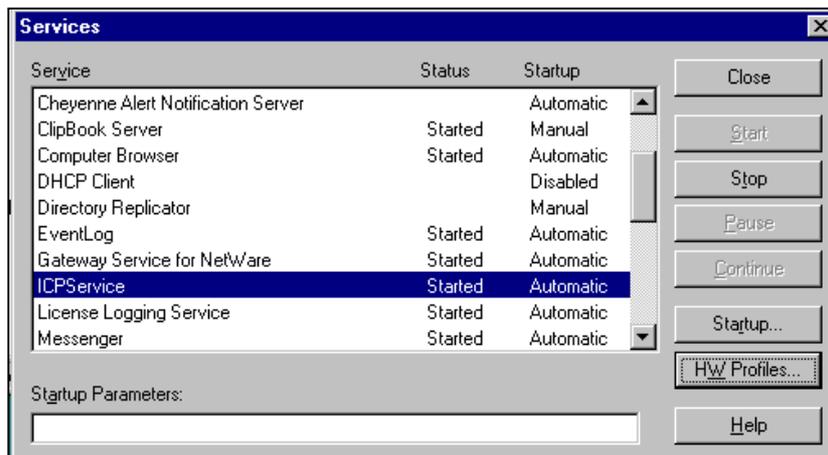
ICP Service	Allows remote access to an ICP Controller in a Windows NT server
ICP CTRLSRV	Allows remote access to an ICP Controller in a Novell Server
ICP Mail	Converts ICP messages into standard mails (for Windows 9x/NT, MAPI format)

It is recommended to install the ICP Service / ICP CTRLSRV and the ICP Mail tool on each server which is equipped with an ICP Controller. Thus, remote access to ICP Controllers in a network can be easily managed from one or several authorized users.

To install the ICP Service under Windows NT, copy ICPSRV.EXE and ICPSRV.CPL into the winnt\system32 directory and load

```
..\ICPSRV -i
```

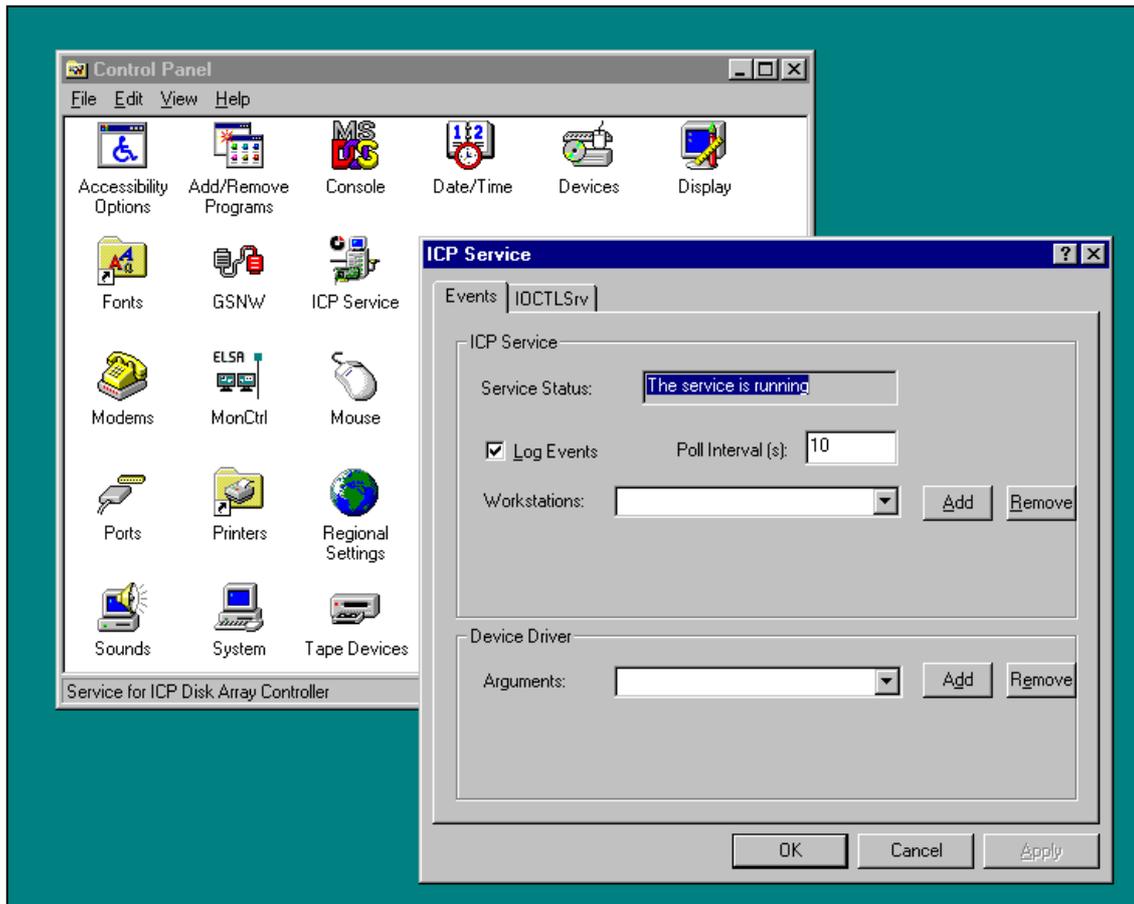
from the command line. This installs the service. In a next step load in the "control panel" the "services" program and activate the ICP Service with the Start button.



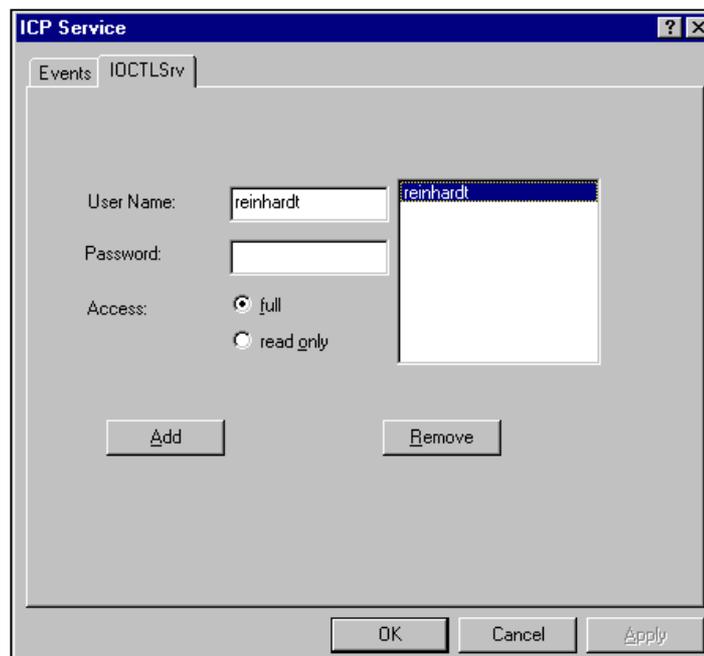
If you want to start the ICP Service during every new boot of the system automatically, click on the Startup button.

To configure the ICP Service double click on the ICP Service icon:

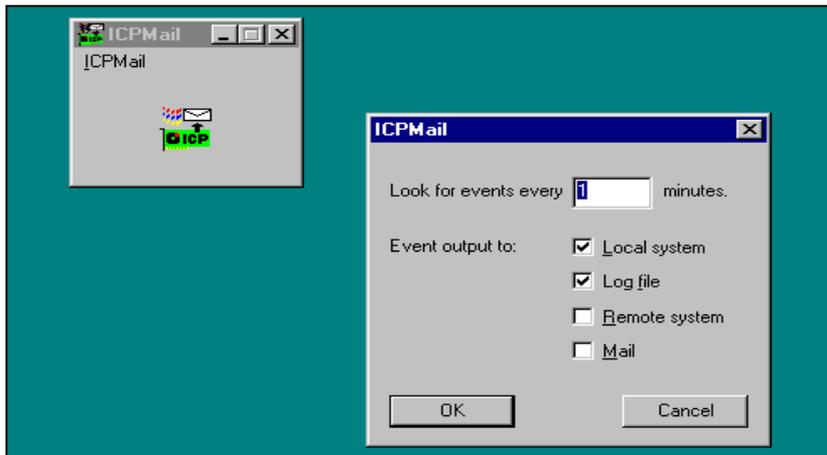
- The Poll-Interval determines the time between two message acquisitions of the ICP Service.
- Under Workstation you can add workstations which receive broadcast messages from the ICP Controller.
- In the device driver section you can add or remove parameters for the gdtx.sys driver (e.g., reserve parameters for raw devices).



In the IOCTLSrv property sheet you can add / remove users which have remote access to the ICP Controller with the ICP RAID Navigator. Passwords are encrypted.



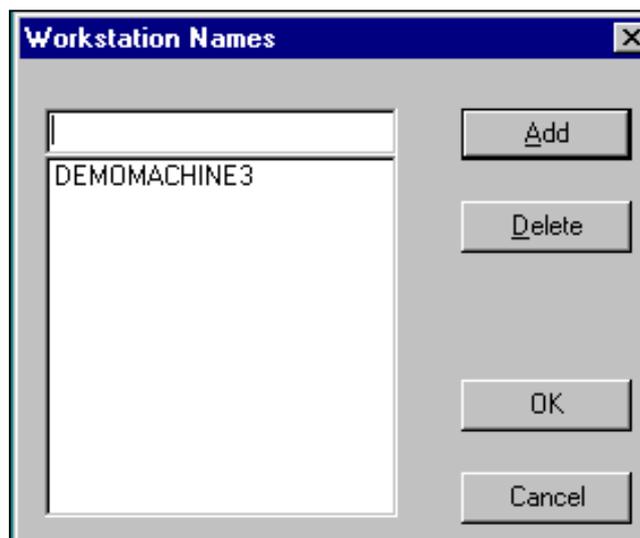
The ICP Mail tool gathers messages from the ICP Service, generates standard mail messages and sends them to pre-defined workstations.
After loading ICPMAIL.EXE and selecting "Settings", you can configure the mailing tool.



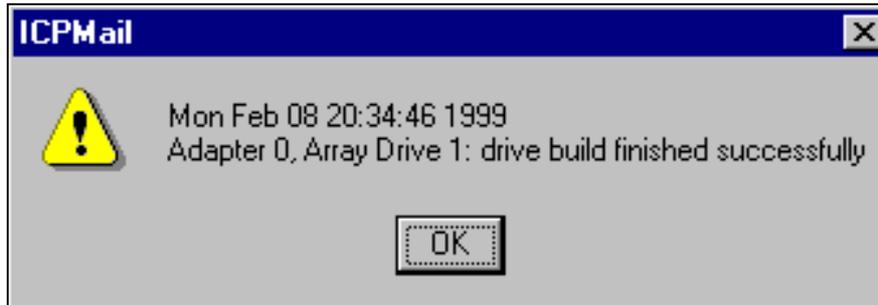
If you select "Local System" all messages are displayed on the server itself.
If you select "Log file" you are asked for a log file path/name. All messages are recorded into this file.



"Remote System" allows you to add workstations to which messages are send.



Following is a typical message generated by ICP Mail.



The "Mail" option allows the interfacing to a standard mailing system (like Microsoft Outlook or Exchange).

Chapter E

Appendix

E. Appendix

E.1 Subsystem Enclosures Manufacturers

Eurologic Systems, Inc., 1300 Massachusetts Avenue, Boxborough, MA 01719, USA
Tel.: 001-978-266-9224
Fax: 001-978-266-9228
Web: www.eurologic.com

Eurologic, Mr. Stefan Kollar
Frühlingstr. 3, 82110 Germering - Germany
Tel.: +89-89445514
Fax: +89-89445522
Web: www.eurologic.com

JMR Electronics GmbH, Mr. Rainer Lehmann
Tel.: +8022-705570
E-Mail: Rainer@jmr-europe.de
Germany

JMR Electronics, Inc., 20400 Plummer Street, Chatsworth, CA 91311, USA
Tel.: 001-818-993-4801
Fax.: 001-818-993-9173
Web: www.jmr.com

Trimm UK, Phoenix Centre, Millenium Way, Nottingham, NG8 6AR
Tel.: 0044-1159519951
Fax: 0044-1159519899
Web: www.trimm.co.uk

E.2 Index

- Acoustical Alarm 68
- Backup 6
- BBS - Mailbox 14
- Cluster Administrator 37
- Cluster Module 13
- CTRLSRV.CFG 43
- CTRLSRV.NLM 43
- DB9 connectors 31
- Deinitialize Disks 62
- Disk Arrays, Build State 51
- Disk Arrays, Error Status 52
- Disk Arrays, Expand State 52
- Disk Arrays, Fail State 51
- Disk Arrays, Idle State 51
- Disk Arrays, RAID 0 45, 66
- Disk Arrays, RAID 1 45
- Disk Arrays, RAID 10 45, 66
- Disk Arrays, RAID 4 45, 66
- Disk Arrays, RAID 5 45, 66, 68
- Disk Arrays, Ready State 51
- Disk Arrays, Rebuild State 52
- dual loop configuration 30
- Failover 32
- FCC Compliance 7
- Fibre Channel Hubs 30
- Flash-RAM 15, 57
- GDT_RPFW, file 57
- Grown Defect List 62
- Host based RAID controllers 12
- ICP RAID Console 42
- ICP System CD 18
- ICP vortex Computersysteme GmbH 14
- ICP vortex Corporation 14
- ICPCON, Grown Defects 49
- ICPCON, Last Status information 49
- ICPCON, Save Information 49
- Interconnects 33
- Internet - E-Mail 14
- IOP (Input/Output Processor), ID 30
- Load Balancing 33
- LVDS/SE Terminator 19
- LVDS/SE Terminator for SE/LVDS Hard
Disks with SCA Connector 19
- Media Interface Adapter, MIA, Order No.
8950 31
- NCPE Protocol 42
- NETBIOS Protocol 42
- Nodes 32
- Notes on Novell Netware Orion Clustering
34
- Novell Netware Orion Clustering 34
- NTFS 32
- NVRAM 33
- Primary Defect List 62
- Private Network 33
- Quorum Resource 32
- Replace Drive with ICPCON 67
- SC Duplex Connector 26
- Shared Storage Bus 32
- Software License Agreement 15
- Static Load Balancing 33
- Swap Host Drives 54, 72
- Synchroner Transfer 61
- Warranty 4