
VA Linux System's RAID Monitor (VARmon)

VARmon provides the ability to quickly and easily determine the status of a large or small RAID system. It also provides some RAID administration utilities for knowledgeable users. The name VARmon itself stands for VA Linux Systems RAID Monitor.

Introduction

VARmon provides a graphical management tool for the Mylex controller products offered by VA Linux Systems. You can use it for:

- viewing the current status of drives and the controllers attached
- rebuilding drives
- marking drives online and offline
- running consistency checks

VARmon also displays statistics from:

- multiple Mylex RAID cards (up to 8)
- multiple physical drives (up to 45 with a maximum of three SCSI channels with 15 physical devices each) per card
- multiple logical drives (up to 32) per card

Note: You must run VARmon with superuser privileges.

Theory

VARmon uses communications read by the driver in the /proc directory. It parses the data into structures which are presented to the user. The information presented to the user is automatically refreshed every 1.5 seconds.

Installation

The VARmon program uses the Linux DAC960 driver that is part of every VA Linux System software load. The VARmon program requires that the DAC driver be at version 2.0.0 or 2.2.0 minimum revision level.

Using VARmon

To start VARmon, follow these steps:

1. Open an terminal with xterm, eterm or a similar window.
 - Windows set in Xwindows for viewing VARmon must be at least 80 by 25 characters wide.
2. Enter this command: **varmon**.

```
=<| VA RAID Monitor (VARMon) |>=
Ver. 1.0.0 by Dragan Stancevic <visitor@valinux.com>

Detected cards:
DAC1164P 3C

Status  Devices  Arrays  Expert
Card:    DAC1164P      PCI Bus:    3
Channels: 3           Device:     8
Memory:  32MB          Function:   0
Firmware: 5.07-0-07    IRQ:       19
-Controller Settings-  PCI Address: 0xF4210000
Controller Q Depth: 128  Mapped to:  0x88000000
Driver Q Depth: 127
Stripe Size: 64KB      -SAF-TE Information-
Segment Size: 8KB      C0 Fan0:ON   Fan1:ON   Fan2:NONE
Max Block/Command: 128 C0 Temp0:39 °C  Temp1:39 °C
Max Scatter/Gather: 33 C0 Power0:ON  Power1:N/A
BIOS Geometry: 255/63
SAF-TE: ON
-General Driver Info-
Author: Leonard N. Zubkoff
Ver.: 2.2.4
Kernel: 2.2.12-29.4smp
Host: colate

<ESC> - Exit | <R> - Refresh | <RETURN> - Lock
```

Figure 1 VARmon Status Menu

All RAID controller cards in the system detected at powerup are shown in the left column. As Figure 2 shows, you can view information about the controller highlighted in the left column by selecting from the tab menu selections along the upper right side of the screen below the credits. Initially, the status menu tab is active and its information is shown in white lettering on a blue background.

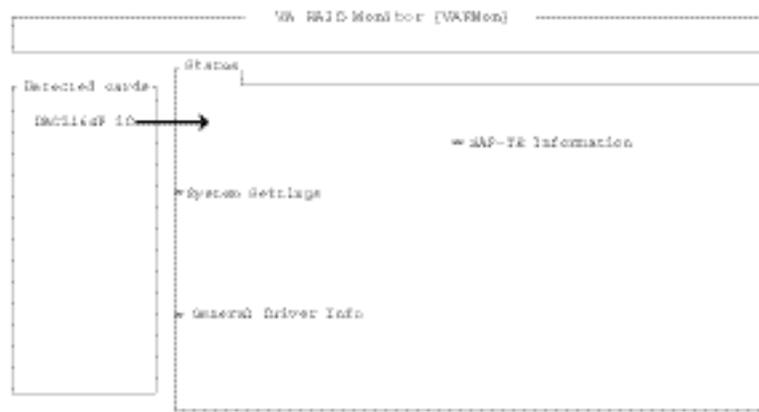


Figure 2 Tab Menu Selections Apply to Highlighted Card

The inactive tab menu selections are displayed in yellow. They are:

- Devices
- Arrays
- Expert

Begin by using the right and left arrow keys to highlight the different tab menus.

The initial status menu presents information about the highlighted RAID card. If your immediate concern is the status of a failing drive, use the right arrow key to highlight the Devices tab, and check the status of the drives attached to that card.

Note: Note that as you move the cursor with the right arrow key the Devices tab menu turns to white.

See the following tables for information key use and visual indicators in VARmon.

Table 1 Visual Indicators

Color or Symbol	Meaning
White Tab	Active menu
Yellow Tab	Inactive menu
Yellow device or card	Array or card in RAID array in critical state. Either a failure has occurred or a rebuild is underway
Red device or card	A drive or card has failed, and the associated data has been lost, if it wasn't backed up.
Green device or card	A drive or card is functioning normally.

Use the following keys for VARmon menu selections:

Table 2 VARmon Key Use

Key	Use
Up and down arrow keys	RAID card selection
Right and left arrow keys	Tab menu selection
<Enter>	Select highlighted choice
<Esc>	Exit current selection
<F5> key	Refresh screen

Status Tab

The status tab selection provides status about the card's properties and settings. In the detected cards column, all the Mylex cards that are detected at powerup are shown.

```

=<| VA RAID Monitor (VARMon) |>=
Ver. 1.0.0 by Dragan Stancevic <visitor@valinux.com>

Status  Devices  Arrays  Expert
-----
Detected cards
DAC1164P 3C

Card: DAC1164P          PCI Bus: 3
Channels: 3            Device: 8
Memory: 32MB           Function: 0
Firmware: 5.07-0-07    IRQ: 19
PCI Address: 0xF4210000
Mapped to: 0x88000000

-Controller Settings-
Controller Q Depth:128
Driver Q Depth: 127
Stripe Size: 64KB
Segment Size: 8KB
Max Block/Command: 128
Max Scatter/Gather:33
BIOS Geometry: 255/63
SAF-TE: ON

-General Driver Info-
Author:Leonard N. Zubkoff
Ver.: 2.2.4
Kernel:2.2.12-29.4smp
Host: colate

SAF-TE Information
C0 Fan0:ON Fan1:ON Fan2:NONE
C0 Temp0:39 °C Temp1:39 °C
C0 Power0:ON Power1:N/A

<ESC> - Exit | <R> - Refresh | <RETURN> - Lock

```

Figure 3 Status Tab

This menu provides information about the controller settings. Most of the information in this screen is gleaned from the driver, however, several items in the Status menu require further explanation:

BIOS Geometry The 255/63 setting indicates that a software geometry fix is in place to support SCSI hard drive that are larger than 8 GB in size.

SAF-TE Information For enclosures that support the SAF-TE (SCSI Accessed Fault Tolerant Enclosure) monitoring protocols such as the FullOn 2x2, temperature readings and fan status from the sensors within each chassis. In the example shown above, the status from a single enclosure is displayed. If multiple chassis are attached, the status from each (up to three) is displayed.

Devices Tab

When you select the Devices tab, you view status information about the SCSI hard drives connected to each SCSI channel. The information is presented in the form:

CH stands for SCSI Channel
ID stands the SCSI ID number of a SCSI hard drive

The status (STAT) codes are:

ONL is online
SBY is standby
DED is dead
WON is write only (this indicates that the drive is in the rebuild process)
Size is the size in GB of the applicable SCSI hard drive

Note: The SCSI hard drive size in GB shown by VARmon will be smaller than that calculated by the SCSI hard drive manufacturer. Manufacturers calculate a GB to be 1,000,000,000 bytes while VARmon calculates a GB to be 2^{30} .

In this menu:

- green indicates that the logical drives defined by the indicated card are functioning correctly
- yellow indicates the logical drives defined by the indicated card are critical
- red indicates the presence of drives defined by the indicated card that are dead

```
Ver. 1.0.0 by Dragan Stancevic <visitor@valinux.com>
```

Status		Devices		Arrays		Expert	
CH:ID	STAT	Size	CH:1	No Devices	CH:2	No Devices	
0:1	ONL	17.1GB					
0:2	ONL	17.1GB					
0:3	ONL	17.1GB					
0:4	ONL	17.1GB					
0:9	FullOn	2x2					

<ESC> - Exit | <R> - Refresh | <RETURN> - Lock

Figure 4 Devices Tab

Figure 4 shows four SCSI hard drives connected to a single Mylex SCSI controller channel.

Note: Figure 4 shows the FullOn 2x2. In the FullOn 2x2, the SAF-TE midplane uses SCSI ID 9, so the STAT and Size categories are not applicable to this listing.

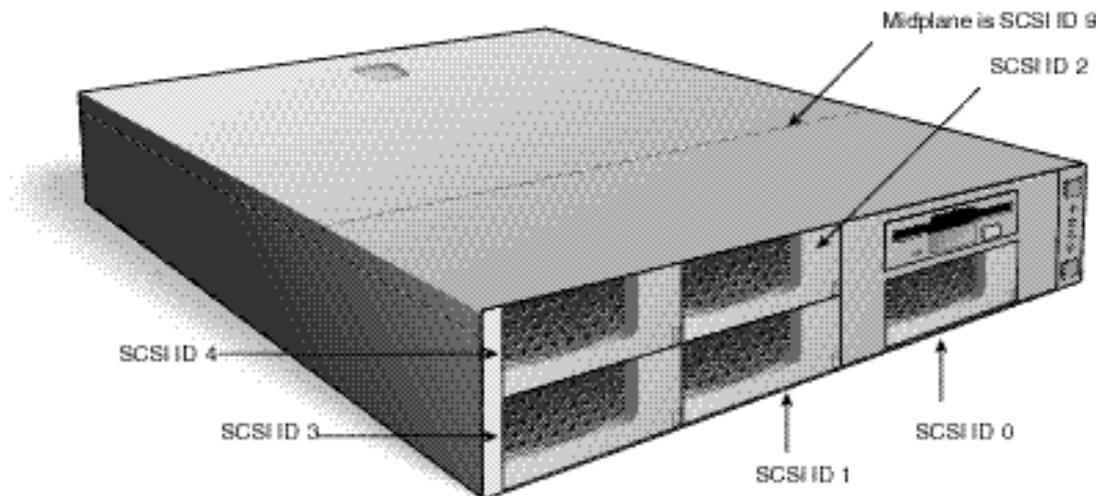


Figure 5 FullOn 2x2's Removable SCSI Hard Drive ID Numbers

Example of Hot-swap Hard Drive Removal

When a hard drive fails, console messages similar to the following appears:

DAC960#X: Physical Drive X:X is now DEAD

If you remove a hot-swap drive, console messages similar to the following appear:

DAC960#X: Physical Drive X:X killed because it was removed

DAC960#X: Physical Drive X:X is now DEAD

The RAID controller detects the drive failure and continues to operate in a degraded mode. If you check the Devices screen, the failed drive is shown in red followed by a *. Also, if you check the Arrays tab, the applicable array might have entered a critical state as shown in Figure 8. If you have the FullOn 2x2, the LED on the hard drive carrier displays solid red as shown in Figure 6.

```

--<I VA RAID Monitor (VARMon) I--
Ver. 1.0.0 by Dragan Stancevic <visitor@valinux.com>

```

Status		Devices	Arrays	Expert
CH:ID	STAT	Size	CH:1 No Devices	CH:2 No Devices
0:1	ONL	17.1GB		
0:2	ONL	17.1GB		
0:3	DED*	17.1GB		
0:4	ONL	17.1GB		
0:9	Fullon	2x2		

<ESC> - Exit | <R> - Refresh | <RETURN> - Lock

Figure 6 Dead Drive

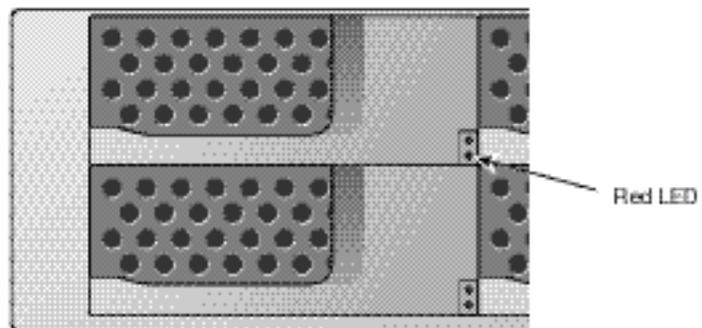


Figure 7 FullOn 2x2 Drive Status Red LED Location

The console displays:

DAC960#X: Logical Drive X (/dev/rd/c0d0) is now critical

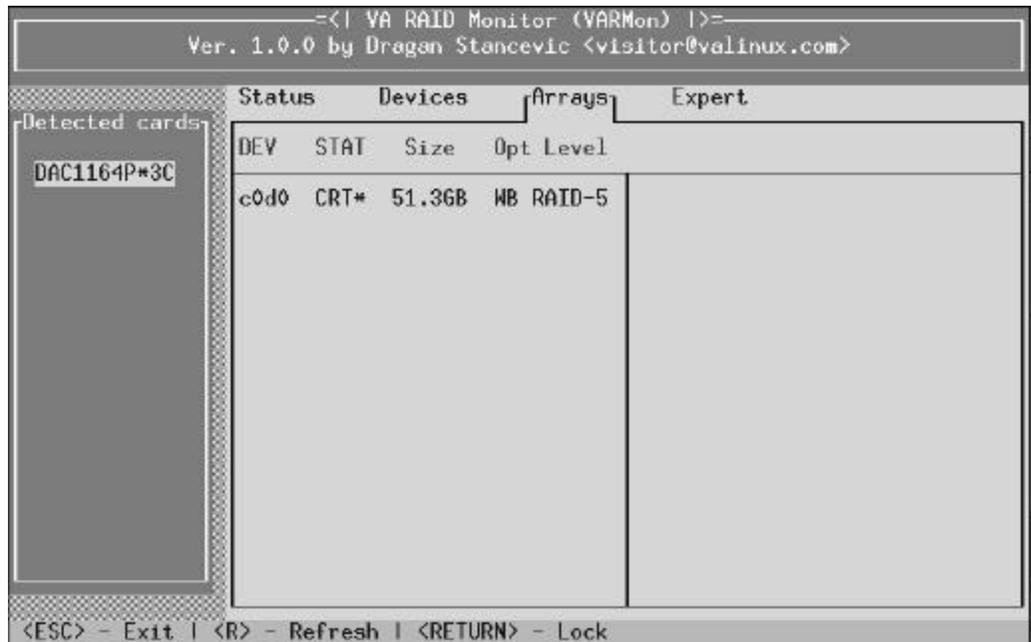


Figure 8 Array In Critical State

After you replace the drive, the SCSI hard drive spins up and the array is automatically rebuilt.

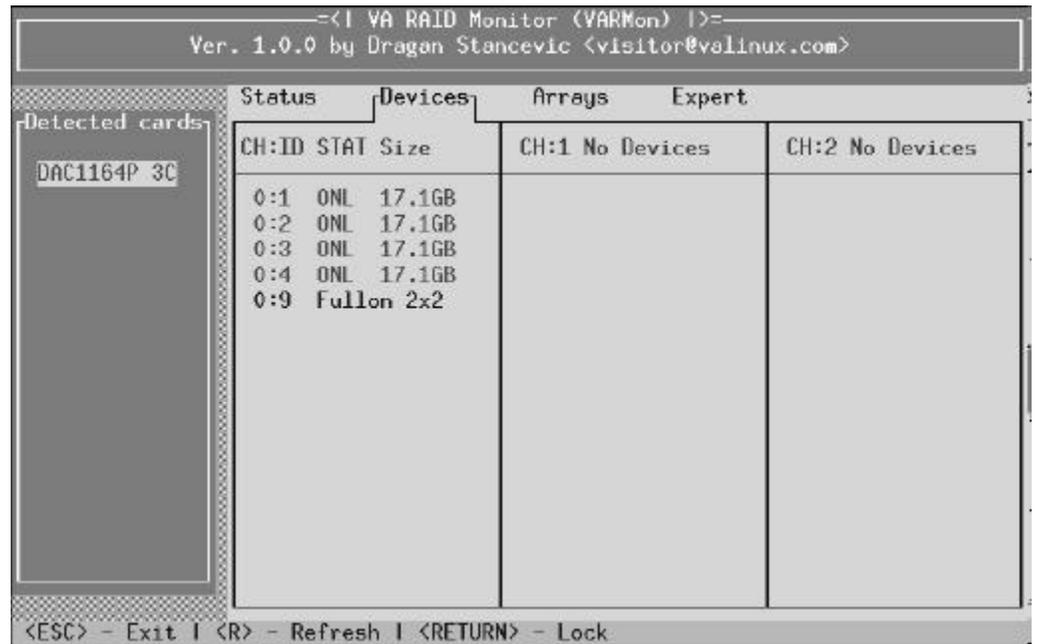


Figure 9 Normal SCSI Drive Display

The LED on the hard drive turns to green during the rebuild.

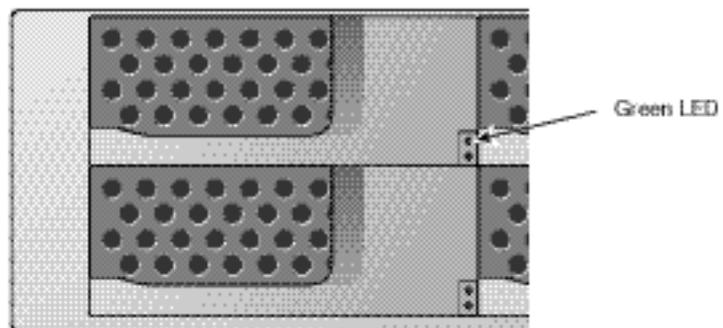


Figure 10 FullOn 2x2 Green LED During Rebuild

The console displays:

DAC960#X: Rebuild in progress: Logical Drive X (dev/rd X% completed)

If you highlight the Expert tab, you can monitor the progress of the array as it rebuilds as shown in Figure 11.

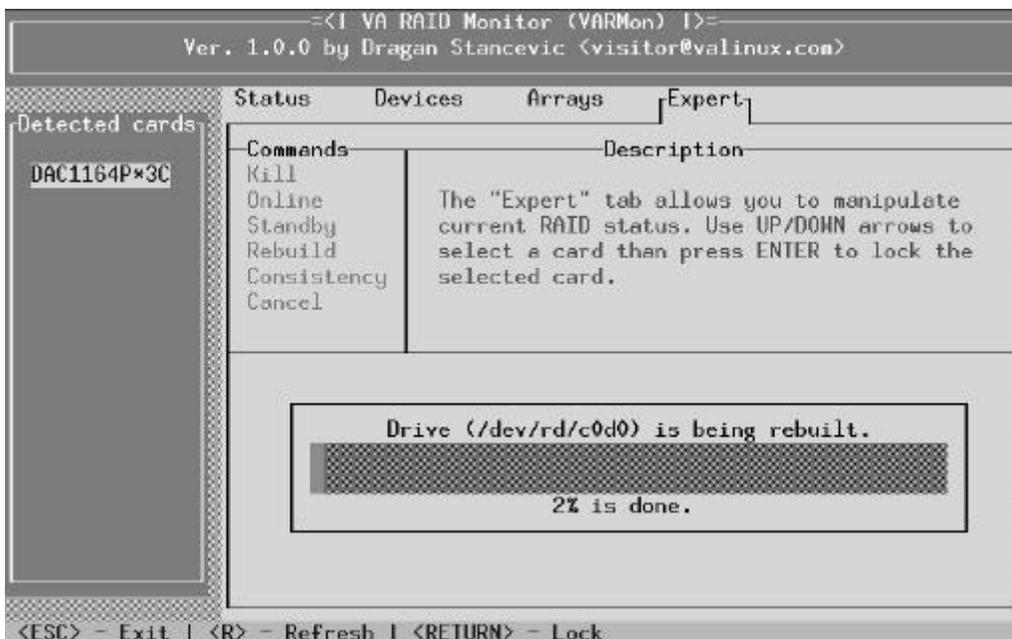


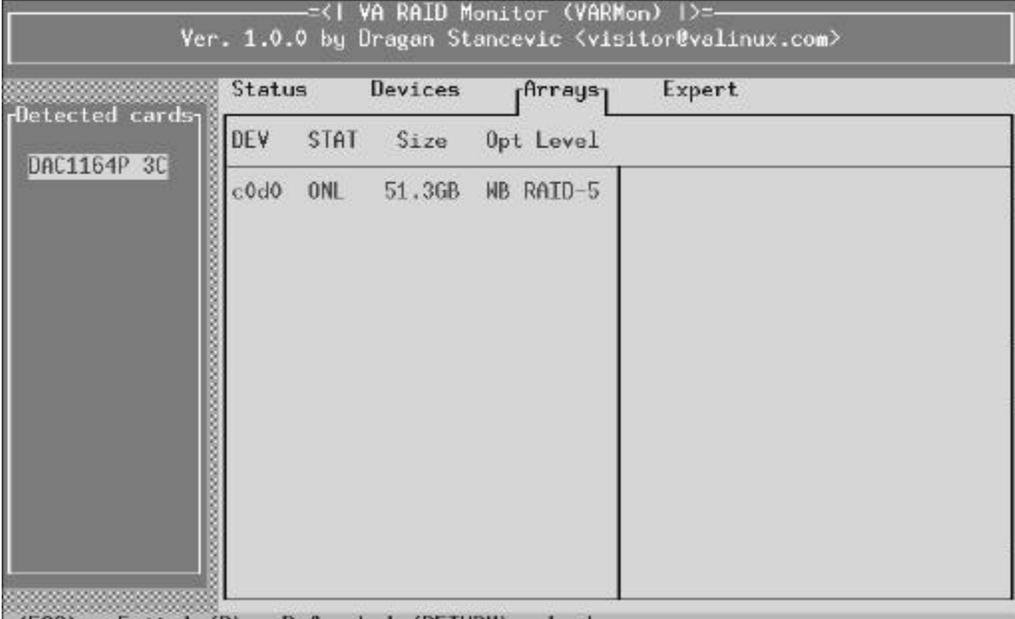
Figure 11 Rebuild Underway

Arrays Tab

The tab entitled “Arrays” shows the status of each RAID array. Figure 12 shows this tab.

In this menu:

- green color indicates that the logical drives (same term as system drives) defined by the indicated card are functioning correctly
- yellow color indicates the logical drives defined by the indicated card are critical
- red color indicates the presence of drives defined by the indicated card that are offline



The screenshot shows the VARmon interface with the 'Arrays' tab selected. The title bar reads '-<I VA RAID Monitor (VARMon) I>-' and the version is 'Ver. 1.0.0 by Dragan Stancevic <visitor@valinux.com>'. The 'Detected cards' list shows 'DAC1164P 3C'. The main table displays the RAID array status:

DEV	STAT	Size	Opt Level
c0d0	ONL	51.3GB	WB RAID-5

Navigation options at the bottom: <ESC> - Exit | <R> - Refresh | <RETURN> - Lock

Figure 12 Arrays Tab

Expert Tab

The expert screen is shown in Figure X.

Warning: As the name of this tab indicates, the options in this menu are for RAID experts. If you make the wrong choices, your RAID array's data will be lost permanently unless you have a working daily backup of the data on removable media in place. If you aren't sure about the options in this section, don't use this menu, go find an expert!

However, this menu is useful to all viewers because it shows whether or not a RAID array being rebuilt, and if so, indicates the level of rebuild completion. As the prompts indicate, select the card that has hard drives connected that you wish to manage and press **enter**.

Use the following keys for menu selection in this menu:

Table 3 Expert Tab Key Use

Key	Use
Up and down arrow keys	Select Kill, Online, Standby, Rebuild, Consistency and Cancel options.
<Return>	Select or "lock" highlighted choice
<Esc>	Exit current selection
<r> key	Refresh screen after DAC driver console messages

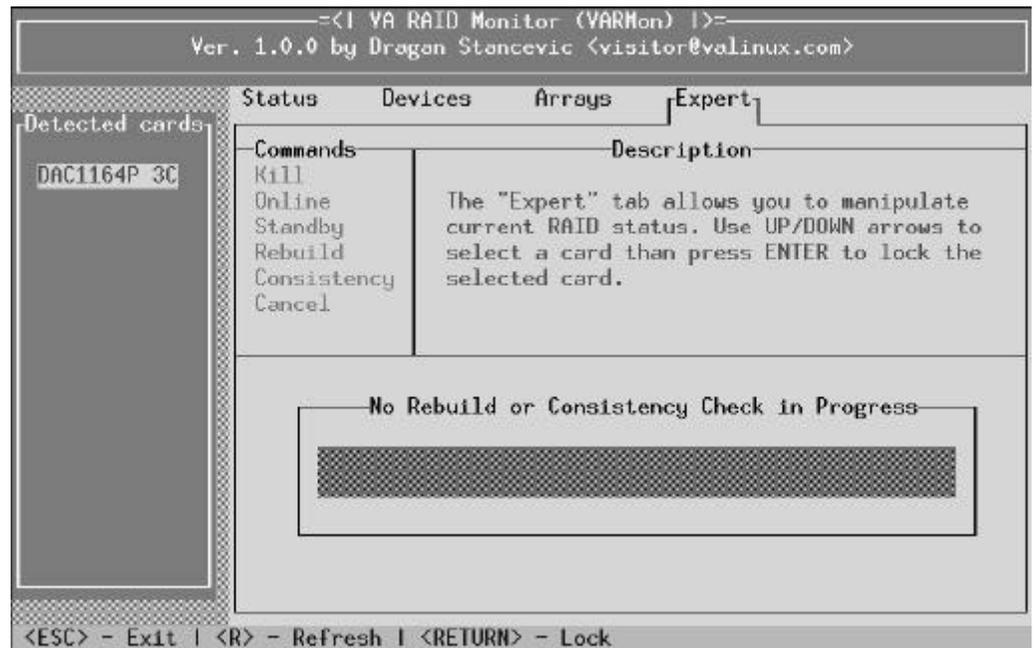


Figure 13 Expert Tab

Using the Kill Command

The kill command is an option in the expert screen. To kill a drive, follow these steps:

1. Select a card, and press **enter**.
2. Select the kill option as shown in Figure 14.

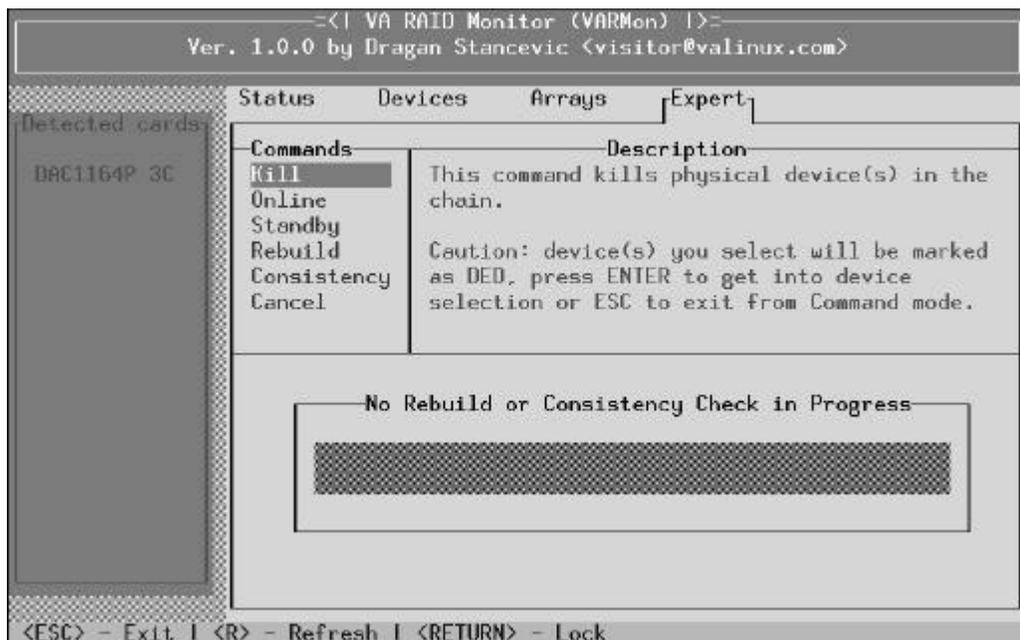


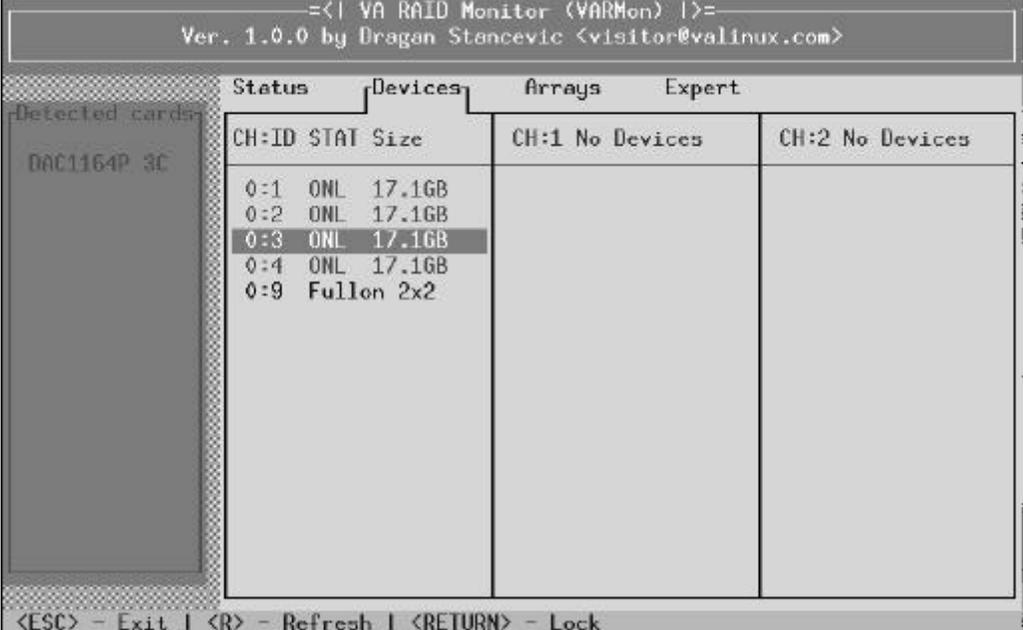
Figure 14 Kill Option in the Expert Menu

3. Select the Devices tab and scroll to the hard drive that you wish to kill with the **up** and **down** keys. Highlight the applicable hard drive. Figure 15 shows the third hard drive on the first SCSI channel selected.
4. Press **enter** to kill the drive.

The console displays a message similar to the following:

DAC960#0 Physical drive X:XX is now DEAD

5. To go back to the expert tab options, and press **esc**.



The screenshot shows the VARmon interface with the following components:

- Header: =(<| VA RAID Monitor (VARMon) |>=
- Version: Ver. 1.0.0 by Dragan Stancevic <visitor@vallinux.com>
- Left Panel: Detected cards: DAC1164P 3C
- Main Table:

Status		Devices	Arrays	Expert
CH:ID	STAT	Size	CH:1 No Devices	CH:2 No Devices
0:1	ONL	17.1GB		
0:2	ONL	17.1GB		
0:3	ONL	17.1GB		
0:4	ONL	17.1GB		
0:9	Fullon	2x2		

Footer: <ESC> - Exit | <R> - Refresh | <RETURN> - Lock

Figure 15 Selecting a Drive

VARmon cautions the user about this choice as shown in Figure 16.

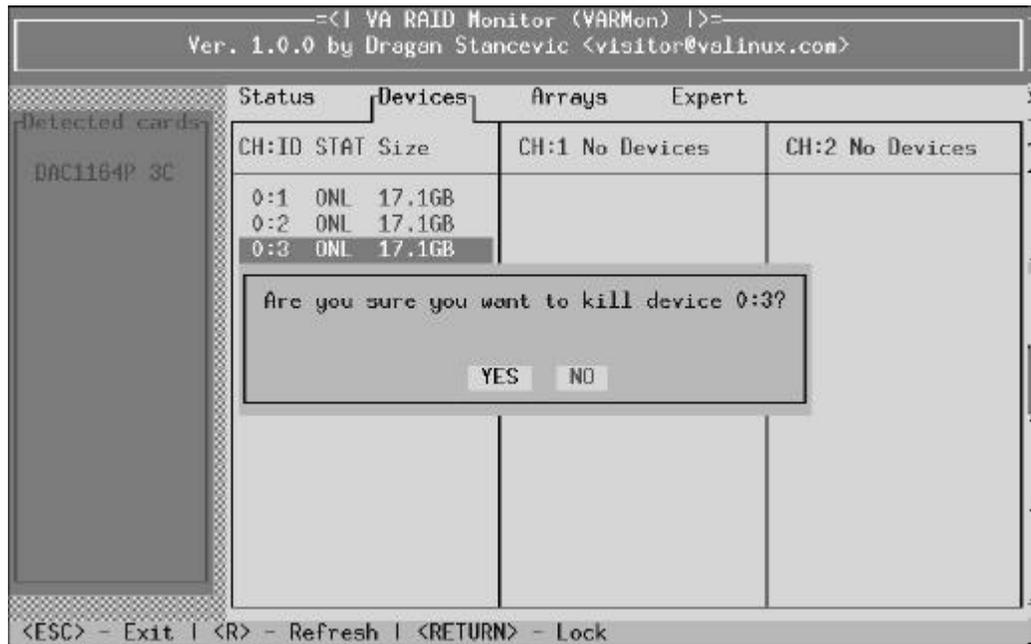


Figure 16 Kill Warning



Figure 17 Notification of Successful Killing of SCSI Hard Drive ID 3

The Devices screen shown in Figure 17 now shows a DED* (dead) drive.

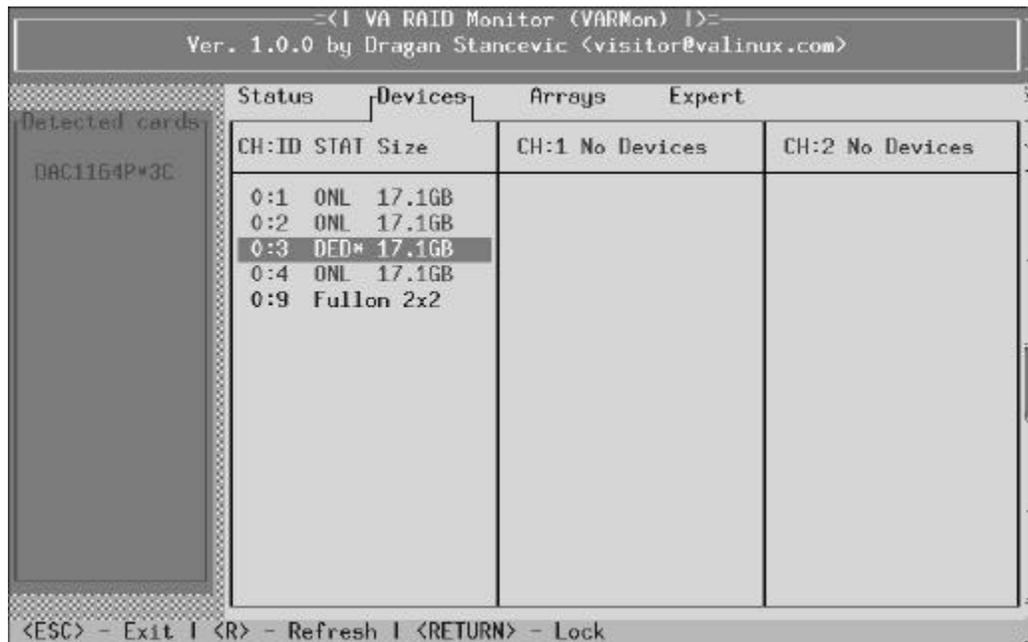


Figure 18 Dead Drive

In the example shown above, the dead drive is swapped, and the rebuild to a spare SCSI hard drive begins as shown in Figure 19. During the rebuild the spare SCSI hard drive is marked as being accessible for writes only (WON) as shown in Figure 20.

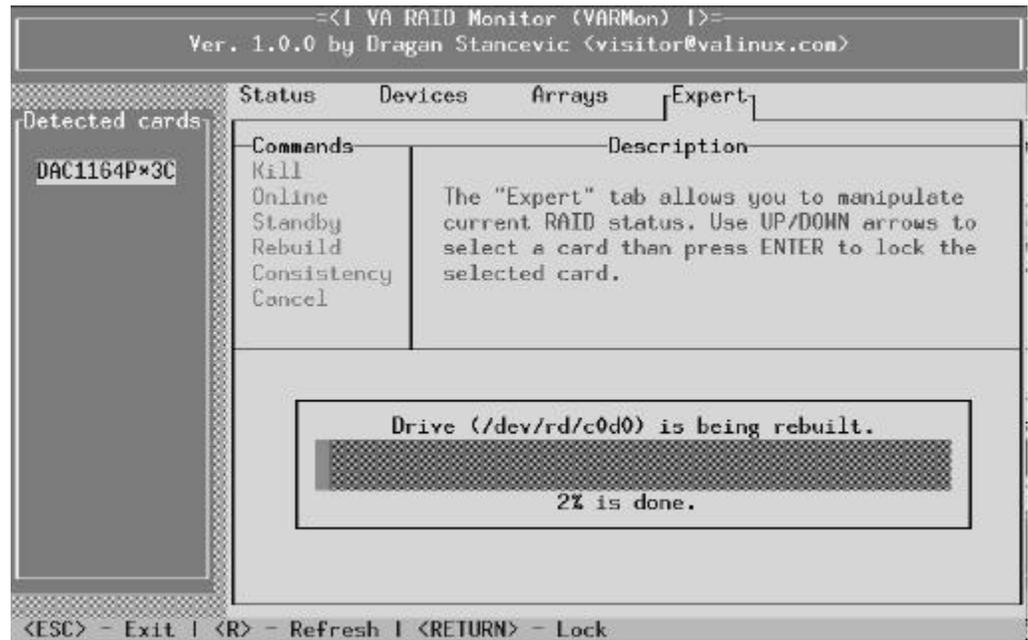


Figure 19 Rebuild Underway

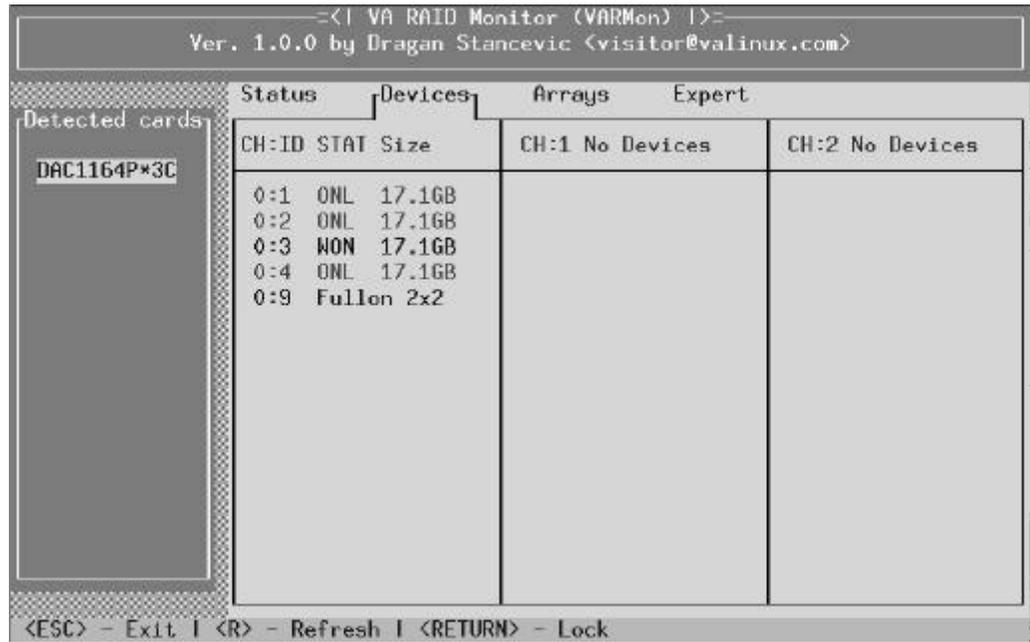


Figure 20 Drive 3 Write Only During Rebuild

Online Selection

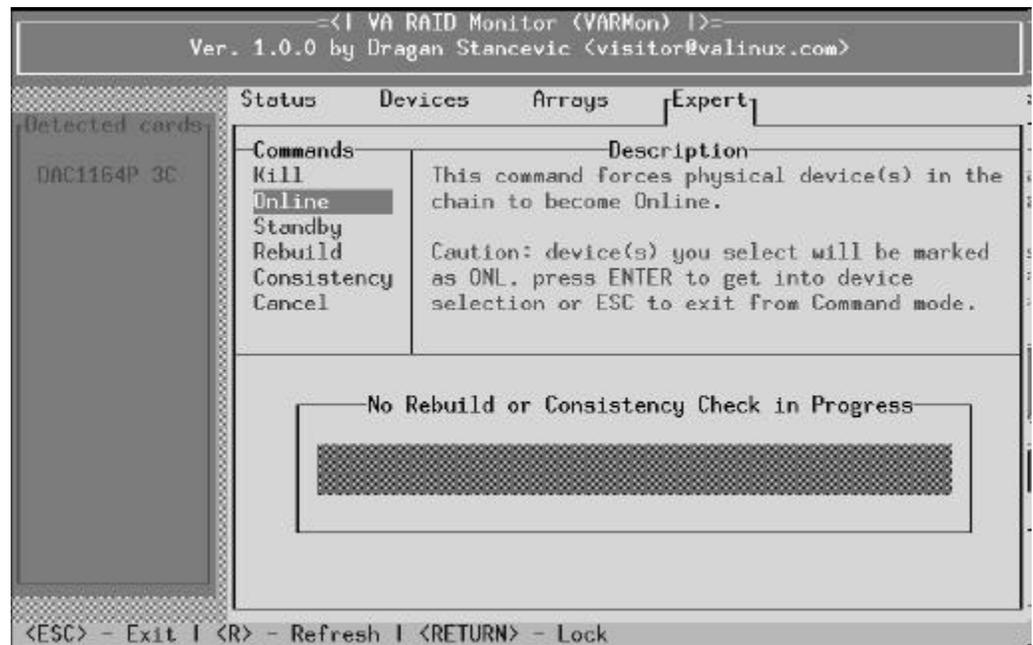


Figure 21 Online Selection from the Expert Menu

The online selection enables a SCSI hard drive.

Standby Option

Configuring a standby SCSI hard drive enables an automatic failover when a SCSI hard drive failure occurs. The standby drive needs to be of the same type as those in the array.

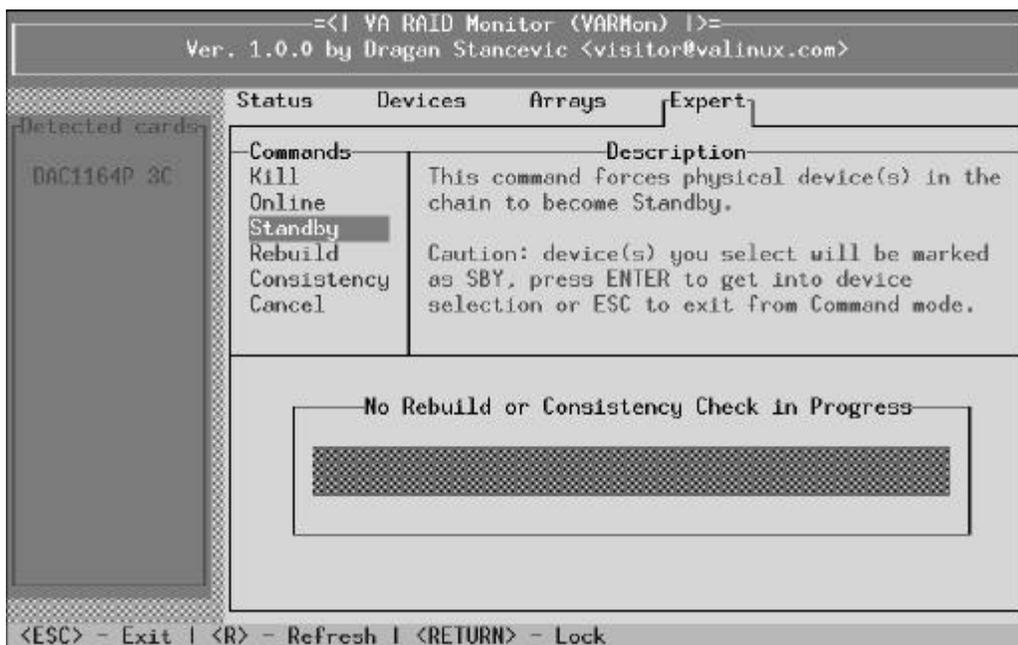


Figure 22 Standby Option from the Expert Menu.

Rebuild Selection

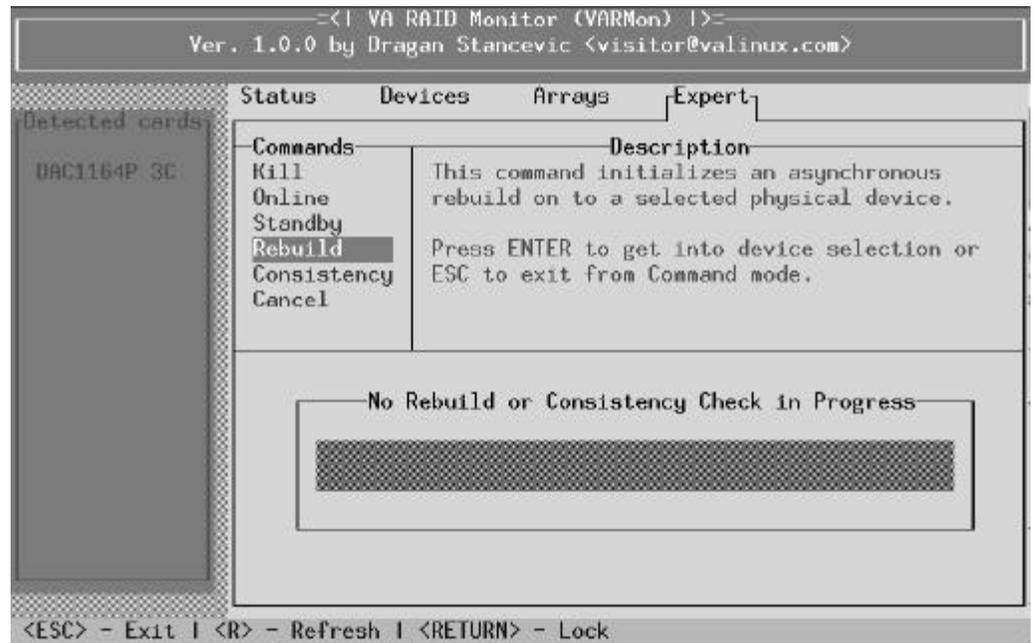


Figure 23 Rebuild Menu from the Expert Menu.

If you have hot-swapped a new SCSI hard drive into the chassis to replace a hard drive failure, you could use the Rebuild selection to rebuild an existing array.

Consistency Selection

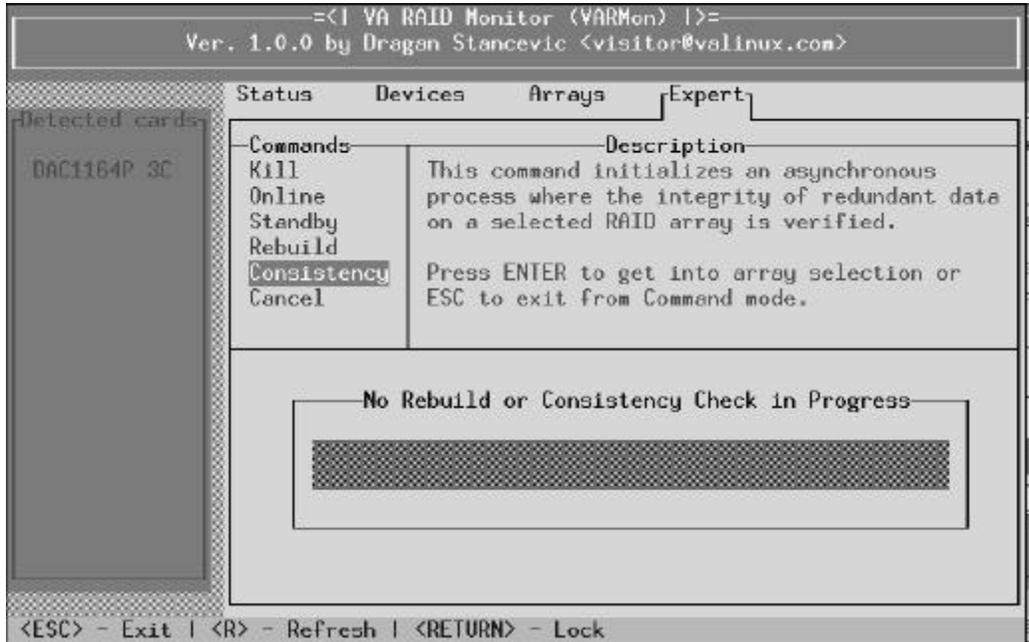


Figure 24 Consistency Option from the Expert Menu.

The consistency option initiates a low level hardware process in which the data on the SCSI hard drives is checked for soundness. For example, if a customer brings a new SCSI hard drive online, then a consistency check provides a method to increase confidence in the integrity of the array.

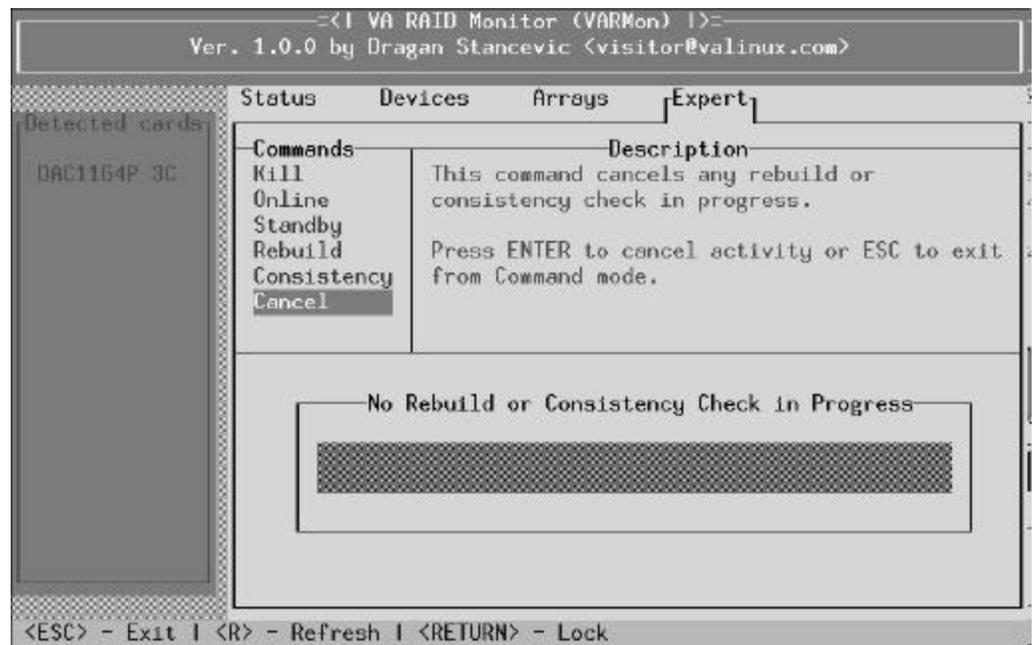
Cancel Selection

Figure 25 Cancel Option from the Expert Menu

This selection is used to stop any rebuilds or consistency checks that might be underway.

