Chapter 25
RA_DISKS

Chapter rewritten by Brian Hailstone, updated by Dave Bazley, with contributions from Richard Penn, Tony Smith, Vic Taylor and others.
Chapter revision information : Rewritten April 1992, updated Sept. 1992 and Sept. 1994

25.1 Confusion surrounding formatting

This section originally by Stefan Fiala...

Clarification of issues surrounding formatting on ALL DSA drives.

25.1.1 A few DSA terms

LBN	Logical Block Number. The place where the users and the operating system store data files and programs. For drives covered in this chapter, 1 logical block = 1 physical sector = 512 data bytes.
RBN	Replacement Block Number. Used for replacement of defective blocks in the LBN area.
RCT	Replacement Control Table. The place where any replaced blocks are kept track of.
FCT	Format Control Table. The place where the media serial number and the location of blocks found bad by the factory are stored. A bit called "fake" is also held here. The FCT is only accessible by the controller.
DBN	Diagnostic Block Number. A place where the controller can perform read/write tests without messing up customer data.
XBN	eXternal Block Number. External blocks are the FCT area.
BBR	Bad Block Replacement. The process of replacing a failing/failed block with another block.
EDC	Error Detecting Code. 16 bit code used to detect data path problems through the controller.
ECC	Error Correcting Code. 170 bit Reed-Solomon code used for error detection and correction.
FE	Forced Error, an error that occurs if reading a block that was written with questionable data during a bad block replacement.

The object of the exercise is to make the LBN area a contiguous entity of good storage blocks. For instance if a customer buys an RA81 consisting of 456 MB, he gets a drive with 456 MB of usable area. Or LBN's 0 to 891,071 (decimal) 891,072 * 512 (bytes/LBN) =456.228 MB. The customer could use this as one file of 891,072 contiguous blocks without worrying about what would happen IF.....

Digital Internal Use Only

The architecture masks media defects from the user. DSA drives have DYNAMIC BBR. In other words, blocks that become unreliable are replaced before they die completely.

There are two slightly different ways of performing BBR.

Controller Initiated BBR. The controller does EVERYTHING, NO operating system involvement at all. The process should be reported to the errorlog. For example, HSC50/70, KDM70 and RQDX1/2/3 use this mechanism.

Host Initiated BBR. The controller notifies the operating system that a BBR is necessary it is then up to the operating system to interact with the controller to perform the replacement. UDA50 and KDA50-Q use this mechanism.

25.1.2 Impact of the FORMATTER

The RCT will over time get more and more entries as more and more blocks are replaced. The increase will be dependent on the inherent reliability of the media. The more Bits Per Inch (hence cheaper/more dense) the more unreliable. RA80's being not very dense only have few entries after say one year while RA82's may have several tens of entries. You will now realise that preservation of the RCT is very important on "used" media.

25.1.2.1 What the two formatter modes do

REFORMAT.	First ZERO the RCT table. Second take the FCT bad-block entries and make corresponding entries in the RCT. Third scan the media and replace any blocks found bad. The FCT is left alone. Any FIELD entries are destroyed. Has minimal impact on brand new (virgin) media. Should not be necessary on "used" media unless to put blocks erroneously replaced back into service [after a fault is CURED].
RECONSTRUCT.	First ZERO the FCT then ZERO the RCT. THIS DESTROYS EVERYTHING. Second, WRITE the FAKE bit to the FCT. Third, scan the media and replace any blocks found bad. ALL bad block entries are DESTROYED. ALSO this disk can now NEVER be reformatted as the FAKE bit is SET. Subsequent requests to re-format will cause a reconstruct. The FAKE bit cannot be altered in the field. The media is one step removed from USELESS.

The method for scanning the media in both cases will normally not find much wrong. In both cases expect a spate of errorlog entries for ECC type errors which should diminish with free space usage and time.

25.1.3 Forced errors

The forced error flag indicates to the host that incorrect data has been correctly written into a block. When an uncorrectable ECC error occurs in a block, the block is re-read a number of times in an attempt to get good data from that block. When the data is read from an uncorrectable failing block and written into the replacement block, the data is a best guess only. To inform to user that the data in the block was at one time uncorrectable, the forced error flag is attached to the block.

The mechanism used to flag a forced error is to complement the EDC character when writing the block. This would appear as an EDC error during EVRLA/ZUDC, with the good data being the inverse of the received data (FE = EDC complemented).

A BBR is requested/instigated on receipt of data that causes the ECC algorithm to exceed a drive specified threshold, **nothing else.** This BBR would also have an errorlog entry for an ECC error (over threshold) and also an errorlog entry for the BBR with a success/fail code.

Digital Internal Use Only

25.2 Testing RA disk drives

-This section originally by Stuart Rance...

There seem to be a number of misunderstandings about how to set about testing a disk drive attached to an HSC, UDA50. Here are some tips that may help you. A disk drive on an HSC or UDA50 can be in any one of three states.

ONLINE means that at least one host has mounted the drive and it is not available for testing. When a drive is online the port select button will be depressed and lit on the selected port. The other port button may be either in or out but only one port button will be lit.

AVAILABLE means that the HSC or UDA50 can see the drive (i.e. that the port button is depressed and the SDI cables connected correctly) but that no CPU is currently accessing the drive. When a drive is available the port button will be depressed but unlit.

The OFFLINE state is shared by offline drives, non-existent drives and drives that are so faulty that the controller cannot see them.

To tell which state a drive is in type a ^C on the HSC console and then type SHOW DISKS. AVAILABLE and ONLINE disks will be shown as being in that state. Other drives will not be shown.

If a drive is ONLINE then simply pressing the port button will not cause it to go offline, the port select button will remain lit and the drive will still be able to respond to commands from its controller. The only way to cause a drive to go offline is for the software to dismount the drive. (Or in some extreme cases to power the drive off). Similarly pushing the write protect button will not change the state of an online drive unless the software wishes to honour the request.

If you want to run diagnostics from an HSC you will need to turn the enable/secure switch on the HSC to the enable position. BEFORE YOU TURN THE ENABLE/SECURE SWITCH you must type a few characters on the console terminal. Otherwise it is possible that a break character stored in the controller could put the HSC into ODT.

If you manage to get the HSC into ODT by accident you will receive one of two prompts, @ or *. To put it back to normal type P to an @ or type ;P to a *.

25.2.1 Online Diagnostics

Table 25–1: Host based diagnostics/utilities

VAX	PDP11	MDM	Comments
EVRLB	ZUDK		Formatter (NOT A DIAGNOSTIC)
EVRLF	ZUDH		Tests 1, 2 and 3 1 = Unibus interrupt and address tests 2 = Execute drive resident tests 3 = Disk function test (R/W to DBN's)
EVRLG	ZDUI		Test 4, disk exerciser
EVRLJ	ZDUJ		Test 5, UDA/KDA subsystem
EVRLK	ZDUL		BBR utility (scrubber)
EVRLL	ZDUM		Disk resident errorlog dump utility
EVRAE			MSCP subsystem exerciser
		NAKDAx	MicroVAX, tests 1-4, format, BRR utility etc

There are NO on line diagnostics for the HSC disk drives, the only way to test these drives under VMS is to use UETP. UETP is a reasonably good confidence test of the drive but it will not usually help you to diagnose a fault to FRU. If you want to use UETP to test just a particular drive then here are some points to watch out for.

Get the customer to either disable disk quota's for the drive or give [1,7] a quota equal to at least half of the free space on the drive. Get the customer to create a [SYSTEST] directory on the drive. Edit the file SYS$TEST:UETINIDEV.DAT to make this drive the only testable device on the system. This file has a format that looks like this (the line END OF UETINIDEV.DAT is essential).

```
DDB  T  HSC001$DUA      (HSC001 DUA drives are testable)
UCB  T  0000            (drive 0 is Testable)
UCB  N  0001            (drive 1 is Not testable)
DDB  T  MFA             (TM78 on this node is testable)
UCB  N  0000            (drive 0 is Not testable)
UCB  T  0003            (drive 3 is testable
END OF UETINIDEV.DAT
```

UETP will only test a particular device if both the DDB line and the UCB line for that particular drive are marked T (i.e. testable). So you should edit the file to make the drive you want to test the only testable device. If UETINIDEV.DAT does not exist the easiest way to create it is to start uetp and then to abort it after the initialisation phase, alternatively you can just create the file yourself with your favourite editor.

Now your UETINIDEV.DAT is ready all you have to do is to log in to the systest account and type RUN UETDISK00 this will prompt for a controller designation to which you reply with the disk name (leaving out the unit number). At this stage if you have got everything right UETP will start testing your drive. If you want to run an extended test then type DEFINE MODE LOOP before you run UETDISK00 this will cause the test to loop continuously until you type a ^C.

Do not abort this test with a ^Y otherwise the temporary files will not be deleted and you may find that the customer is short of thousands of blocks of disk space.

25.2.2 In Line Diagnostics

The HSC has its own diagnostics which can be used for diagnosing disk faults to FRU. These are ILDISK and ILEXER. They can both be run from the HSC console terminal.

BEFORE YOU RUN EITHER OF THESE DIAGNOSTICS YOU MUST MAKE SURE THAT THE DRIVE IS NOT ONLINE.

Type a ^C on the HSC console and then type SHOW DISKS. Check that the drive you want to test is not online. If it is then ask the customer to take it offline before you continue.

25.2.2.1 Running ILEXER

ILEXER is a multi drive exerciser that will simultaneously test HSC disks and tapes. It will only run on disks that are in the available state. Type RUN ILEXER (for HSC50 type RUN DD1:ILEXER) to the HSC> prompt, the program will then prompt you for a drive unit number, type Dnnn to test disk drive number nnn. ILEXER will then prompt you for other bits of information, such as whether you want to write to the customer area of the pack, eventually it will ask ANOTHER DRIVE(Y/N) []? if you answer yes then you can input details for another drive. If you answer no then you will get another set of prompts including RUN TIME IN MINUTES, ENABLE SOFT ERROR REPORTS etc. Once ILEXER has started it can be quite difficult to abort. There are times during the testing when it will ignore ^Y's. Just keep typing ^Y and eventually you will get through.

25.2.2.2 Running ILDISK

ILDISK is a disk drive diagnostic that will run on any disk drive that is not online to a host. That is it will run on an available or on an unknown drive. It causes the selected disk drive to execute its resident micro-diagnostics.

Digital Internal Use Only

To run it type RUN ILDISK (for HSC50 type RUN DD1:ILDISK). ILDISK will then prompt for a drive unit number. Type Dnnn where nnn is the unit number. If the HSC cannot see the specified drive it will prompt you for the requestor number and port number, you will have to find these for yourself either from the site management guide or by following cables.

ILDISK will then ask if you want to run a single drive diagnostic. If you answer no then it will run the entire suite of drive diagnostics otherwise you will have to enter the number of a particular routine. These drive routines are drive specific and can be found in the orange pocket guides for each drive. If you want to abort the test then type a ^Y, at certain times it can take up to two minutes before the diagnostic is able to halt, just keep typing ^Y's.

25.2.3 HSC utilites

The HSC also has some utilites the may be used to as an aid to diagnosing RA disk problems.

25.2.3.1 VERIFY

The VERIFY utility allows you to check the integrity of the architectural structure of a disk and ensure that it conforms to DSDF. VERIFY **only reads** the disk and therefore does not destroy user data. It does not perform BBR but will report errors.

25.2.3.2 DKUTIL

DKUTIL can be used to display disk structures and disk blocks.

One of the most useful things DKUTIL can do for you is type the error silo (internal errorlog) from the drive. Do not use this command on an RA60.

```
HSC>RUN DKUTIL (or RUN DD1:DKUTIL)

DKUTIL>GET Dxxx  (xxx = disk drive unit number)
DKUTIL>DISPLAY ERRORS  (can be shortened to DI E )
```

DKUTIL can also display the RCT and FCT. This will show which blocks have been replaced.

```
DKUTIL>GET Dxxx
DKUTIL>DISPLAY FCT
DKUTIL>DISPLAY RCT
```

25.2.4 Running Drive Resident Micro-diagnostics

These diagnostics are different on each of the drives supported, there are some pre-requisites that apply to all of the drives.

The drive must not be ONLINE to any host. This state can be identified by typing SHOW DISKS on the HSC console. If the drive is shown as being online then get the customer to dismount it from all systems that have it mounted. When the disk is offline the port select buttons should go out.

The drive must not be AVAILABLE to the HSC. In practice what this means is that when the drive is not online you should release both port select buttons.

A useful tip for testing the RA's is to release both port select buttons and the run button, (first check that the drive is not online). Then spin the drive up and it will loop in self test on spin up until a port select button is pressed in.

For further information on drive resident diagnostics see the relevant pocket maintenance guide.

Digital Internal Use Only

25.3 Drive General Information

25.3.1 RA Drive Specifications

Figure 25–1: Disk heads layout

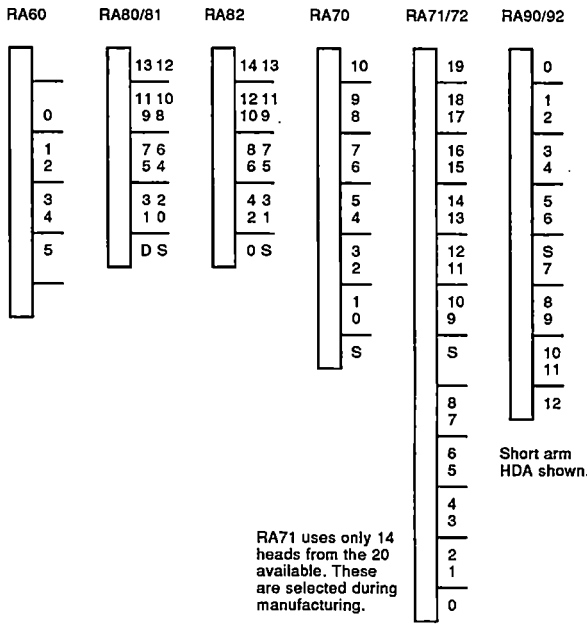

Digital Internal Use Only

Table 25–2: RA disk drive data

	RA60	RA70	RA71	RA72	RA73
Capacity Mbytes	205	280	700	1000	2007
Physical cylinders	1600	1517	1929	1929	2681
No. of data heads	6	11	14	20	21
LBN's per track	42	33	51	51	70
User LBN total	401184	547041	1367310	1953300	3920490
Max transfer rate MHz	15.8	11.6	16.0	16.0	22.1
ECC threshold	4	6	6	6	6

	RA80	RA81	RA82	RA90	RA92
Capacity Mbytes	122	456	623	1216	1505
Physical cylinders	559	1258	1435	2656	3108
No. of data heads	14	14	15	13	13
LBN's per track	31	51	57	69	73
User LBN total	238700	891072	1216665	2376153	2940951
Max transfer rate MHz	9.7	17.4	19.2	22.2	22.2
ECC threshold	2	6	6	6	6

25.3.2 HDA Handling

The way in which HDAs are handled by Engineers and others can have a significant effect on their future reliability. Remember HDAs are fragile, expensive and often in short supply. All HDAs should be treated carefully, the following points apply to the older (RA8x) drives:

1. To close the top chassis only a light force should be required. If not, adjust the latch. Always spin down the drive before closing the chassis.
2. To slide the drive in or out, spin it down but leave the power on.
3. Always use the RUN/STOP switch to spin a drive up or down. **Never** just power it off. If it doesn't spin down immediately, wait... up to 5 minutes!
4. Never wrench the drive out on it's runners against the half-way locks.
5. Ensure the HDA is not resting on any cable clamps. Check drives at PMs, remove the HDA if necessary.
6. Follow this procedure to avoid "head scrape on stationary disk":

   To *Lock* the heads:
   With the drive spun up, release RUN switch and wait a few seconds for heads to return to HOME, then lock with lever.

   To *Unlock* the heads:
   Ensure the power is on when you move the lever.

   In both these cases power is applied to the positioner to hold the heads in place whilst the lever is moved.

Digital Internal Use Only

25.3.3 Head/Cyl from LBN conversion

How to convert LBN's to cylinder/head using a calculator, with a choice of two different methods.

25.3.3.1 Method 1

- Cylinder = LBN/X (discard the fraction)
- Head = (LBN - (cylinder * X))/Y (discard the fraction)

Drive	X	Y
RA70	363	33
RA71	714	51
RA72	1020	51
RA80	434	31
RA81	714	51
RA82	855	57
RA90	897	69
RA92	949	73

LBN conversion example.

```
    LBN number 321456 on an RA82
LBN/X                      = 321456 / 855           = 375.9719 = 375
(LBN - (cylinder * X))/Y = (321456 - (375*855)/57 =  14.5789 = 14
This gives cylinder 375, head 14
```

25.3.3.2 Method 2

1. Divide LBN by X, then discard the fraction
2. Divide the result by Y, the whole number is the cylinder, drop the whole number
3. Multiply the fraction by Y, result is the head number

Drive	X	Y
RA60	168	6
RA70	33	11
RA71	51	14
RA72	51	20
RA80	31	14
RA81	51	14
RA82	57	15
RA90	69	13
RA92	73	13

LBN conversion example.

```
    LBN number 321456 on an RA82
1.  LBN/X       = 321456 / 57 = 5639.578947
2.  5639/Y      =   5639 / 15 =  375.933333
3.  0.933333*Y  = 0.933333*15 =   13.999999  (14)
This gives cylinder 375, head 14
```


Digital Internal Use Only

25.4 Drive specific information

25.4.1 RA60

25.4.1.1 Measuring RA60 runout

Typically a drive suffering from runout problems may exhibit error codes 26,27,28,29,2C,2F. To measure the runout of a RA60 pack/spindle use the following procedure.

- Set up your scope as follows :-
  1. Set channel 1 to 1V/div
  2. Set timebase to 5ms
  3. Set trigger to External positive
- Put channel 1 probe on Test Point 2-2 of the drive logic module. (See Figure 25–2)
- Put trigger probe on Test Point 2-7 of the drive logic module. (See Figure 25–2)
- Connect HHT or a terminal to the RA60 diagnostic port.
- Enter the following commands at the RA60> prompt and after each command measure the peak to peak runout.
  1. RA60> 05 399 0 ;select Head 0
  2. RA60> 05 399 2 ;select Head 2
  3. RA60> 05 399 4 ;select Head 4
- Remove the pack and rotate the spindle 90 degrees and carry out steps 1-3 again. Do this twice so effectively you are measuring at three points of the compass.

If any of the three heads at any of the three points displays a runout of more than 1.2 Volts peak to peak the pack or spindle is bad.

Figure 25–2: RA60 Drive Logic Module

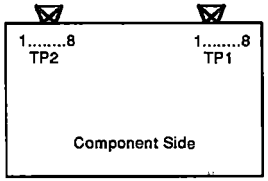

25.4.1.2 RA60 Disk Well Glue Problem

The glue weld joint which attaches the air inlet flange to the disk well assembly can develop cracks on older drives. These cracks can cause an air leak which ultimately can be the cause of head crashes.

To inspect the weld joint for signs of cracks do the following :-

1. Spin drive down, remove pack and power off drive.
2. Remove front bezel.
3. Remove filter clamp and absolute filter.

Digital Internal Use Only

4. Looking from the inside of the disk well, GENTLY press the air inlet from inside the disk well down. Check for any occurrence of a gap between disk well and glued on air inlet flange.
5. Do the same as in 4 looking at the bottom of the air inlet flange.

If cracks are detected the disk well will need to be replaced, the part number is 70-18480-01.

25.4.2 RA70

25.4.2.1 RA70 spin down FCO RA70X-I001

RA70 drives, with an ECM below Rev-J6 (V79 ucode), which have an HDA with a newer motor, can log errors at spindown. These errors do not affect drive performance, they just fill up the errorlog. FCO RA70X-I001 is a microcode upgrade (to V81) which fixes the time-out problem which causes these errors. It should only be applied if the HDA is replaced.

25.4.2.2 RA70 microcode FCO, RA70X-I002

All RA70's should have microcode revision V83 (minimum). FCO RA70X-I002 brings the drive up to V83. FCO kit part number is EQ-01619-02.

25.4.2.3 New RA70 ECM has no ground spring

The RA70 ECM module 70-22494-01, Rev. H4 and above, has no ground spring. The ECM is not DOA because the spring is not there, it's just not needed anymore.

25.4.3 RA71/72

Table 25–3: Operating system minimum version requirements for RA71/72

VMS	V5.4-2[1]
VAXsimPLUS	V1.6
Ultrix-32	V4.2
VAXELN	V4.2

[1]Error Log Formatter (ERF) version 5.4-2 must be upgraded to version 5.4-2 (0001) to support RA71/72.

Table 25–4: Controller minimum revision requirements for RA71/72

HSC40/60/70/90	CRONIC V600	
HSC50	CRONIC V410	
HSC	K.SI requestor SW version 12	
HSC	K.SDI requestor SW version 39/40	
KDM70	SW version 3.0	HW version 17
KDA50	SW version 8	HW version 4
KDB50	SW version 20	HW version 28
UDA50	SW version 6	HW version 0

Digital Internal Use Only

Table 25–5: RA71/72 spare parts

Part no.	Description	RA71/72 spare parts
70-28492-01	HDA RA71	Includes frame assembly
70-28492-02	HDA RA72	Includes frame assembly
54-20826-01	ECM	Electronic Control Assembly
70-29408-01		Skid plate assembly

25.4.4 RA73

Table 25–6: Operating system minimum version requirements for RA73

VMS	V5.5-2
VAXsimPLUS	V2.0
Ultrix-32	V4.3
VAXELN	V4.3

Table 25–7: Controller minimum revision requirements for RA73

HSC40/60/70/90	CRONIC V650	
HSC50	CRONIC V410	
KDM70	SW version 3.0	HW version 17
KDA50	SW version 8	HW version 4
KDB50	SW version 20	HW version 28
UDA50	SW version 6	HW version 0

Table 25–8: RA73 spare parts

Part no.	Description	RA73 spare parts
70-28699-01	HDA RA73	Includes frame assembly
54-21396-01	ECM	Electronic Control Assembly
70-29408-02		Skid plate assembly

Digital Internal Use Only

25.4.5 RA80/81/82

25.4.5.1 RA81 Error code 39 determination

The following procedure should be used when an RA81 is giving error code 39's. It is a procedure to determine if the HDA is the cause of the problem or the electronics. Error code 39 means OFF TRACK , so this procedure measures the runout of the HDA and the servo offset when on track. Of course, if you can determine from the errorlog that all the errors are on one head this procedure is not necessary.

When carrying out the following procedure the signal you will be looking at will be made up of an AC and a DC component, the AC component or ripple should not exceed 0.5V peak to peak,the DC component may be positive or negative but should not exceed 2.5V. If either of the components exceeds the above measurement on any head at any cylinder then an error code 39 would be generated.

- Deselect both port buttons.
- Connect a HHT or terminal to the RA81.
- Connect Channel 1 of your scope to E110 pin 14 on the servo module. (See diagram for location of E110)
- Set Channel 1 of your scope to 1V/division dc coupled.
- At the DIAG> prompt type RUN SEEK
- Now you will be asked to input a cylinder and Group/Head address.
- When the seek has completed look at your scope to see if the signal falls within the above specification.
- Check every head (0-13) and any areas that the errorlog points to.

If the procedure above does not show any head to be out of spec then swap the Servo module and then the microprocessor before swapping the HDA.

Figure 25–3: RA81 Servo module

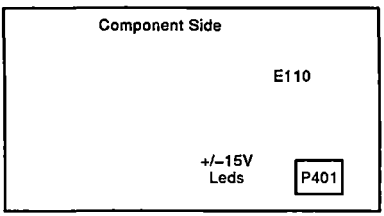

E110 is the sixth chip down in the second row from the right.

25.4.5.2 Problem with RA82 rev H2 Hybrid module

The rev H2 hybrid module for RA82's which fixes a problem with DSM-11 and various port access problems has a minor bug.

If you push the run button and then push it a second time before the drive comes ready you will get F8 errors in the internal error silo. You will not get a fault light or an errorlog entry.

If you see these F8 errors on a RA82 with a rev H2 hybrid module and the drive is otherwise functioning correctly you can safely ignore them.

Digital Internal Use Only

25.4.5.3 Error 2B's on RA82's with new rev KDB50

This problem only affects a small proportion of RA82s. **Do not** replace hardware unless you see the symptoms described and don't forget there are numerous other causes of 2B errors.

After installing the FCO KDB50-R002 (takes microcode to 19) you may see a high incidence of SDI Command Timeout errors, Event code 2B, on RA82's. These errors are not present with the older KDB50 microcode.

If you notice 2B errors after installing the FCO or after replacing the T1002 module with a revision "Mx", or later, check the RA82 Hybrid module. If the Hybrid is not revision "H" or higher replace the Hybrid with a revision "H".

Note: In concept the above scenario could apply to the KDA50 with code revisions 6 or 7 also, but this has not been seen to date.

25.4.5.4 RA8x Preventative Maintenance

There is an RA8x Preventive Maintenance Procedure on fiche, EK-ORA8X-PM-02. Any PM work should take place during, or after a fault call, and the PM procedure should be used as a guide to what to do.

Logistics hold PM kits :-

- RA81 rev 7 70-27494-01
- RA81 rev 8 70-27493-01
- RA82 70-27493-01

They contain the belts, ground springs etc. but not the PSU fans. For RA82 and rev 8 RA81's there is a new cable spring assembly and a shorter drive belt hence the different PM kit no. Read note 7 of PM procedure.

25.4.5.5 Replacing HDAs

When replacing RA8x HDA's always try to ensure that the new HDA has an environment that is going to support it through its life. It is no good putting a HDA into a four year old disk drive without changing the PSU fans and the drive belt, for even if the HDA lasts another 4 years it is almost certain the fans and belt will not.

The only two parameters we can affect in the life of a HDA are its temperature and the stress applied to the spindle. So at every HDA replacement change the PSU fans, clean the front filter, change the drive belt.

25.4.5.6 SAFETY HAZARD on SA482

Several SAFETY hazards exist while installing SA482s and H9646 cabinets.

- DE-skidding is a three man job. Don't attempt it without help.
- Manual EK-H9646-UG-002 Chapter 2, section 2.2.1 has the correct procedure.
- Check the pallet ramps before use for :-
  - Cracks more than 25% of the ramp.
  - Knots or knot holes for more than 50% the width.
  - Loose, missing or broken ramp side rails.
  - Loose, missing or bent metal hardware.
- If the ramp is damaged DON'T use it.

ED's Note: the above applies to installing any cabinet, be careful boys!

Digital Internal Use Only

25.4.5.7 Potential RA82 Power Harness Problem and LED codes 25/30

It was found that some harness wires can get damaged by the threads of one of the spindle motor bolts. The damage can cause intermittent LED codes 25 and 30. Symptoms noted are:

- Drive will spin down, with NO front panel lights lit, resembling a power loss condition.
- The drive may go through <MOUNT VERIFY> stages and possibly even recover; if unsuccessful, Led codes FA and 01 can get reported by the error log.

A new (shorter) belt, which was introduced as part of the RA8X PM PROCEDURE, is contributing to this problem by moving the spindle motor forward. An excess service loop in the harness after the servo FRU has been changed is also reckoned to be a factor.

Closely inspect the power harness for wear of the 'white webbing' and/or rub marks on the front (near the top) of the motor. If the insulation is not worn, reposition the cable as outlined below. If the insulation is worn the harness should be insulated, and repositioned, or replaced (PN# 70-21456-01).

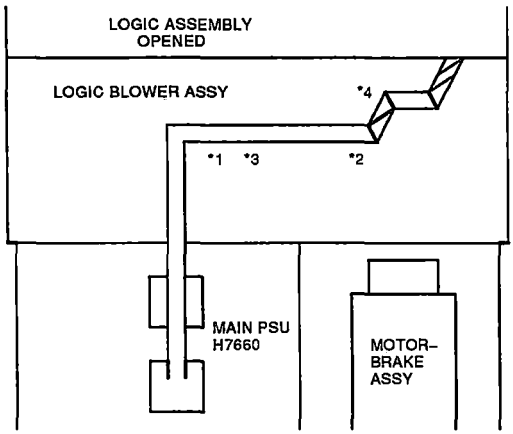

The two looms of the harness are tied to the Logic Fan Assembly at positions *1, *2, and *3. Position *4 is at the right side of the fan assembly, to the rear of the logic chassis, and is used to secure the servo cable (black).

Using a medium size tie wrap, loosely tie the two power harness cables to the same point at position *4. This will provide sufficient clearance between the harness and motor brake assembly.

Digital Internal Use Only

25.4.6 RA90/92

25.4.6.1 RA90 "haze" problem

When RA90s with HDAs prior to rev N12 are spun or powered down, they can fail to spin up correctly. This is the so called "Haze" problem, and is caused by the heads sticking to the platters. This problem is fixed by V27 microcode which unsticks (?) the heads before trying to rotate the disk.

25.4.6.2 Long and Short arm HDAs + ECMs

There are three types of HDA and two types of ECM which can be fitted to RA90s and 92s:

Description	Part No.	Comment
RA90 Long Arm HDA	70-22951-01	Original RA90 HDA
RA90 Short Arm HDA	70-27268-01	Later RA90 HDA (Head arm is shorter)
RA92 HDA	70-27492-01	Only RA92 HDA
ECM Module	70-22942-01	Original ECM
ECM Module	70-22942-02	Later ECM

V25 microcode is the minimum version to support the newer parts. However, V27 is now required for **all** drives so this should be no problem.

25.4.6.3 RA9x microcode FCO, RA9x-F005

All RA90's and RA92's should be upgraded to microcode revision V27. The microcode is available as an FCO. The FCO kit part number is EQ-01643-01, and it contains a microcode cartridge (70-27950-03).

25.4.6.4 RA9x power supply FCO, RA9x-0006

The RA90/92 power supplies manufactured by CEAG have a high risk of failure. Therefore an FCO has been put in place to replace these PSU's. Only CEAG PSU's are affected, ZENITH PSU's are ok. Replace only the CEAG PSU's at revision level "N" and below. The affected serial numbers are as follows;

- RA90 drives from KB00008570 to KB00024239
- RA92 drives below KB00005716

Any drive which had its PSU replaced during a fault call may also have a CEAG power supply that needs replacing.

The FCO part number is EQ-01633-01 and it contains a 30-23506-01 power supply.

Make sure the new PSU is set for 240v before you install it.

Digital Internal Use Only

25.5 Controller and System Information

25.5.1 UDA50 - Unibus Disk Adapter

25.5.1.1 Delay Jumper Setting

A lot of documentation describing this jumper is incorrect. The UDA50 is supplied with the jumper set for a burst latency (time between NPRs) of 0 uS. This should be OK for most systems. It should only be changed if **other** devices suffer from missed transfers. See below for jumper positioning.

Don't confuse burst latency with burst length. This is a software parameter governing the number of tranfers "hogged" per NPR. It should be left alone.

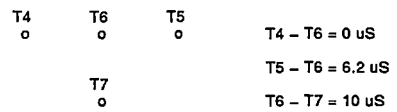

25.5.2 KDA50 - QBus Disk Adapter

25.5.2.1 Delay Jumper Setting

Same as UDA50, see above.

25.5.2.2 Two sorts of SDI module

The M7165, KDA50 SDI module comes in two flavours that look very different. They are however, functionally the same and both work fine with any M7164 (Qbus half of KDA50). The newer module , etch rev E, draws 3.5 amps less of +5 volts.

25.5.3 KDM70 - XMI Disk/Tape controller

25.5.3.1 Overview

The KDM70 controller provides eight SI ports and two data channels. The KDM70 can support up to eight DSA devices, with data transfers occurring on any two ports simultaneously. It will support both tape and disk, with tape support restricted to two tape formatters.

The controller consists of two XMI modules, the policy processor (T2022), and the SI interface (T2023).

25.5.3.2 VMS Issues

The minimum release of VMS with full support for the KDM70 is 5.3-1.

When adding a new device to the XMI the VMS Autosizer will assign the device mnemonic (eg. PUA, PUB) based on the XMI node number. The autosizer configures the XMI nodes from the lowest node number to the highest. Since the BI node adaptor is always the highest XMI node, the KDM70s get configured before the KDBs. In systems with KDBs that are adding KDM70s, the KDB controller mnemonic will change when a KDM70 is installed.eg DUA0 becomes DUB0. This rule applys for tapes as well whether or not a tape drive is attached to the KDM70, MUA0 will become MUB0 and so on including the console TK.

VMS does NOT support dual-pathing between KDM70 and HSCs. Local adapters and HSCs quickly run into unit naming conflicts in VMS. Hence we support local to local failover, and HSC to HSC failover with no mixing.

25.5.3.3 Installation

The KDM70 requires a minimum external SI cable length of 12 feet. Cables less than 12 feet may result in SI pulse errors (Event Code = 10B).

The KDM70 internal SI cables connect directly to the XMI back- plane. If not connected correctly, or to the proper backplane slot, module selftest, or port failures may occur. Refer to the KDM70 Users Guide for assistance in connecting these cables.

The KDM70 will pass self-test without internal SI (back plane to bulkhead) cables attached.

Cables on wrong module cause core test failure, LED = F

The KDM70 internal bus is used for communication between T2022 and T2023 module. If this cable is not connected properly, module self-test will fail. Fails board 2 HIB test, LED = 5

25.5.3.4 KDM70 as the Boot device

For a 6xxx CPU ROM and EEPROM list versus XMI options please read Datadoc page 25-40.

25.5.3.5 MINIMUM KDM70 Revision Level

To determine the KDM70 revision level, examine the XMI device type register:

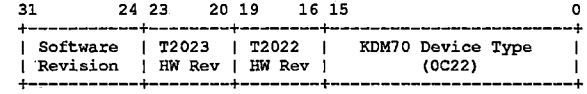

KDM Software...Software readable rev = 2 ... ECO rev 136

T2023 Module...Software readable rev = 1 ... ECO rev C

T2022 Module...Software readable rev = 1 ... ECO rev E

25.5.3.6 KDM70 SOFTWARE Releases

KDM70's ordered prior to December 90 will not have had the software ordered on the BOM, bill of materials which means you have to register the installation to get future release's. Eventually this will become automatic as per Chronic software on HSC's. Any installation's not receiving software at installation time needs to be registered, this is best done by notifying your FCO co-ordinator to track and pass on to the SSB/SPA loop.

25.5.3.7 KDM70 DIAGNOSTICS

The KDM70 provides diagnostics and utilities to assist in sub- system verification and fault isolation. Several programs are provided as part of the KDM70 code image located on the policy processor module (T2022).

Table 25–9: KDM70 RESIDENT DIAGNOSTICS

Program Name	Description
ILDEVO	In-line Device Operations Test
ILEXER	In-line Multi-drive Exerciser
DKUTIL	DSA Disk Utility
FORMAT	DSA Disk Formatter Utility
VERIFY	DSA Disk Verify Utility

Digital Internal Use Only

25.5.3.8 VAX DIAGNOSTICS

VDS based diagnostics will normally be located in the SYS$MAINTENANCE area of the system disk.

Table 25–10: VAX DIAGNOSTICS

Program Name	Description
EVRLM.EXE	Code Update Utility
EVRLN.EXE	Standalone DUP Driver
EVRLN.KDM	Diagnostic Key
EVRLJ.EXE	Disk Subsystem Exerciser
EVRAE.EXE	On-Line Disk Exerciser

25.5.3.9 Invoking KDM70 Controller-resident Programs On Line from VMS

You cannot run on-line diagnostics, exercisers, and utilities without first running EVRLN.KDM. It is important that you follow this procedure.

$ RUN SYS$SYSTEM:SYSGEN

SYSGEN> CONNECT FYA0/NOADAPTER

SYSGEN> EXIT

$ SET DEFAULT SYS$MAINTENANCE

$SET HOST/DUP/SERVER=DUP/LOAD=EVRLN.KDM PUA0/DEVICE

$SET HOST/DUP/SERVER=DUP/TASK=ILEXER PUA0/DEVICE

25.5.3.10 Invoking KDM70 Controller-resident Programs Off-Line from VDS

To invoke EVRLN, boot the VAX/DS in standalone mode. At the VAX/DS prompt,attach and select the KDM70 controller. The following is an example of this procedure.

DS> ATTACH KDM70 HUB DUx N BR

DS> SELECT DUx

DS> RUN EVRLN

EVRLN> RUNL ILEXER

25.5.3.11 EVRLN Overview

EVRLN is a standalone DUP control program that runs under the VAX diagnostic supervisor. EVRLN enables the LOCAL and SUPPLIED DUP diagnostics and utilities to run when on-line testing is not possible.EVRLN must be invoked prior to running in-line diagnostics on the KDM70 controller. Multiple controllers must be tested individually.

Table 25–11: EVRLN Commands

Command	Description
HELP	Prints information on EVRLN commands
CAT	Prints a catalog of SUPPLIED and LOCAL DUP programs
INFO	Prints information about the selected controller

Digital Internal Use Only

Table 25–11 (Cont.): EVRLN Commands

Command	Description
RUNS	Executes a SUPPLIED DUP program
RUNL	Executes a LOCAL DUP program
EXIT	Terminates EVRLN and returns control to VAX/DS

25.5.4 VMS 5.5 - New Errorlog entry

A feature of VMS v5.5 is an entry into Errorlog.sys of ERL$LOGSTATUS messages that seemingly contain very little in the way of verbose information as to what is wrong. New functionality now reports blocks, with the forced error flag set, at the same time increasing the error count for the device.

Two labels in the errorlog entry, contain information about what is going in on:

- CRDP$L_MEDIA is the logical block at the start of the transfer
- CDRP$Q_ISOB is the exit status for this transfer

Use the DCL command EXIT %X****** to decipher.

```
$
$ EXIT %X2144
%SYSTEM-F-FORCEDERROR, forced error flagged in last sector read
$
```

25.5.5 DSAERR and FSTERR

DSAERR and FSTERR are both software tools that can be installed under VMS and used to simplify the analysis of DSA (and other storage) errors. They typically use the information in the errorlog and format it into a concise report.

DSAERR is the older utility and only deals with disk errors.

FSTERR is newer and deals with disk and tape errors. The latest version, V4.0 also supports SCSI and StorageWorks.

They can both be generated from WELSWS. A 1600 BPI tape containing FSTERR is kept in the Datadoc Media Cabinet at Welwyn.

Digital Internal Use Only

F A C T F L A S H

Options Affected: 6410 thru 6460
Submitted By: Dave Bazley
Date: 2-JAN-1991
Filing Instructions: Keep at end of Chapter 25, Calypso

VMS V5.4-1 and 64xx Systems

VMS version 5.4-1 will not work on 6000-400 series systems. It causes the system to crash during booting, with an "INVEXCEPTN, Exception while above ASTDEL or on interrupt stack" bugcheck.

A patched version, VMS 5.4-1A, should be used instead.

This information is explained in the Covering Letter for Owners of 6000-400 Series Systems which is included in the VMS 5.4-1 kit.

Digital Internal Use Only

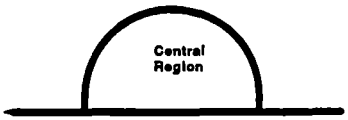

F A C T F L A S H

Options Affected: 6000 series systems.
Submitted By: Pete Griffin
Date: 14-DEC-1990
Filing Instructions: Keep at the end of the new Chapter 25.

Round-up of current issues.

- I've made quite a few alterations to Chapter 25, the 6x00 chapter.
  Mostly it is stuff for the new 6500.
  It would be worth scanning if you are already trained on 6400/6300/6200, to get some understanding of the differences (you will not get "conversion" training).
  I would suggest you pay particular attention to:-
  New approach to updating CPU EEPROMs, **UPDATE must not be used anymore.**
  New ROMs needed to support new XMI devices and Infoserver.
  "Platform" differences
  Upgrade paths
  New PSU and LEDs.
  New FRUs - some are downward compatible.
  New BBU LEDs, fan and switches.
  Booting via the Infoserver
  CIXCD microcode procedures.
- **UETP LOAD phase fails** on systems with VECTOR processors.
  This is fixed by VMS 5.4-1.
- Recently-built **XBIB (T1043), causing errors** (INT60, BI ADAPTER interupt [NAK to MULTI-RESPONDER CMD RECVD] and MEMORY CONTROLLER ERRORs).
  Caused by some non-AMD chips at E18, fixed by T1043 revision K (was J), E18 was 23-007L5-00 becomes 23-050L5-00.
- **6500 Diagnostic Supervisor will not boot** (EMSAA V14.0-561), if there is an I/O device in slot 3.
- 6500 system **must have a console terminal**, else SFT fails.
  A CPU in slot 1 will fail Test 3 of ST1.
  CPUs in any other slot will pass but then hang.

Digital Internal Use Only

- **CIXCD must be set for 10 tick** Quiet Slot Delta time (not 7)

  The "tick time" used to depend on the number of nodes.

  1-5 nodes, set 7 tick on all nodes,
  6-15 nodes, set 10 tick on all nodes,
  16 plus nodes, set 10 tick on all nodes.

  CIBCxs are set up using backplane jumpers E09-E39, E10-E40, E11-E41
  CI750, CI780, and HSCs fitted with L0118s, are set up using L0118 switches.

  16 Node, 7 tick; SW3-1 OFF, SW3-2 OFF, SW3-3 OFF, SW3-4 OFF.
  16 Node, 10 tick; SW3-1 OFF, SW3-2 OFF, SW3-3 OFF, SW3-4 ON.
  32 Node, 7 tick; SW3-1 ON, SW3-2 OFF, SW3-3 OFF, SW3-4 OFF.
  32 Node, 10 tick; SW3-1 ON, SW3-2 OFF, SW3-3 OFF, SW3-4 ON.

  CI750, CI780, and HSCs fitted with L0100 Rev E2, set up using L0100 swtchs.

  7 tick, SW3-1 OFF, SW3-2 OFF.
  10 tick, SW3-1 ON, SW3-2 ON.

  CI750, CI780, and HSCs fitted with L0100 Rev E1, set up using L0100 links.

  7 tick, jumper E177p11 to E177p12.
  10 tick, jumper E177p9 to E177p10.

  CI750, CI780, and HSCs fitted with L0100 Rev D, are 7 tick only (need to fit a L0118).

- **Standalone Backup has a timing problem** configuring HSC disks.
  Symptom is 50 messages saying "%PAA0, Port is re-initializing"
  Fixed by slowing down CPU by disabling the cache.

  1. Method 1 for DEC released media

     Do a conversational boot by setting R5 D0 as a 1.
     SYSBOOT> SET USERD2 3
     SYSBOOT> CONTINUE

  2. Method 2 for @STABACKIT built media

     Boot Standalone Backup and wait for it to ask the time.
     ^P
     >>> DEPOSIT/I 7F 0
     >>> DEPOSIT/I 72 6
     >>> CONTINUE

- VMS 5.3 gives **Fatal Bug Check** - INVEXCEPTION when booting a 6410 as a new cluster member.
  This is due to the new root having DUMPBUG =1, VAXCLUSTER =2, but no SYS-DUMP.DMP.
  Fix by SYSGEN> CREATE SYS$SYSDEVICE:

- 6300 **CACHE TAG PEs** - this should be a soft error, but there is some evidence pointing to dodgy components causing not only this error but also Fatal Bug Checks (INVEXCEPTION etc) due to a cache coherency problem.
  So if you see Cache Tag PEs, it is best to replace the CPU module.

- FCO to **replace all 6400s** with rev L modules is "paused" whilst the cleaning process for the AC terminator chips is devised.

- Can **only boot VDS** thru the CIXCD if you define the HSC in <D15:D08>, <D07:D00> does not work?

- VMS or Standalone Backup, **cannot configure the CIBCA** properly unless it was booted via the CIBCA.
  The problem occurs if a DEBNx or TK is in the BI slot just before the CIBCA, or in the last BI slot of XMI node E, or the first slot of XMI node D.
  The symptom is PAA0, Port timeouts.

Digital Internal Use Only