

EHAS 2900: Volumes Note 1

The following new facilities are available in Volumes version 17A onwards. They are described in detail in the attached note.

- * Exclusive as well as inclusive user lists for BACKUP, ARCHIVE, EXPORT, etc. (see Section 3.2)
- * SUSPEND BACKUP or ARCHIVE or EXPORT (see Section 3.2)
- * File system re-creation (see Section 6)

(File system re-creation is only partially operational as yet. A secure tape is being written nightly, and RECREATE will go through all the motions described herein. However, backup entries are not yet being inserted in users' #ARCHs, so the processing of the latter yields an empty control file; thus no REPLACE requests are generated.)

C.D. McArthur

EMAS 2900 VOLUMES NOTE

The Facilities Provided by the Volumes Process

No: 1
Date: 16/3/79

[This note supersedes all previous Volumes Notes]

Contents

1. Introduction
2. VOLUMS services initiated by users' Directors
3. VOLUMS services initiated by operator
 - 3.1 Management of private devices and volumes
 - 3.2 Management of backup/archive store
 - 3.3 Monitoring
4. Management of backup, archive, export and secure tape lists
5. File system re-creation
6. EMAS 2900 standard tape format

1. Introduction

The Volumes process VOLUMS is privileged (ACR 5) and is started as process number 2. It resides on a fixed site on System discs, although it can be started as any file in the file system.

VOLUMS has a file index in each file system, in which it keeps lists of spool requests, tape lists, lists of files awaiting backup, archive, or export, and log files. (Process 1, DIRECT, also uses these indices, having none of its own).

The services provided by VOLUMS include the management of private devices and volumes and the maintenance of the backup/archive store.

2. VOLUMS Services initiated by users' Directors

The activity numbers of these services lie in the range 64-127. Those currently defined are as follows:

Activity_67: Notification of a user process stopping.

Request: integer DEST, SRCE, P1, P2, P3, P4, P5, P6

Action: Any volume requests outstanding for this user are cancelled, and any volumes allocated to him released.

Reply: None

Activity_68: Request for the allocation of a volume.

Request: integer DEST, SRCE, IDENT, TYPE, PERM, string (7) DSN, c
integer P6

TYPE = 4 - tape (currently the only type defined)

PERM = 1 - read
 = 2 - write
 = 3 - read/write

DSN = identifier of required volume

Action: If the volume is not currently loaded the operators are requested to load it with the appropriate permission; if it is currently loaded with the wrong permission they are requested to unload it and reload it with the required permission. These requests may be refused by the operators. If (or when) the volume is correctly loaded, it is allocated to the requesting user.

Reply: integer DEST, SRCE, IDENT, FLAG, SNO, string(7) DSN, c
integer P6

FLAG = 0 - success
 # 0 - failure

SNO = service number of successfully allocated volume

Activity_69: Request to release a previously allocated volume.

Request: integer DEST, SRCE, IDENT, TYPE, SNO, string (7) DSN, c
 integer P6

 TYPE = 4 - tape
 SNO = service number of volume
 DSN = identifier of volume

Action: The volume is de-allocated from the requesting user, and
 either allocated to another user or unloaded.

Reply: integer DEST, SRCE, IDENT, FLAG, P3, P4, P5, P6

 FLAG = 0 - success
 # 0 - failure

Note: This activity is automatically invoked by a call on activity 67
 ("user process stopping").

Activity_70: Enter a spool request of type "restore" into request list.

Request: integer DEST, SRCE, string (6) TAPENAME, USERNAME, c
 byteinteger FSYS, TYPE, integer CHAPTER, IDENT

 TAPENAME = identifier of source volume
 USERNAME = identifier of requesting user
 FSYS = file system number of requesting user
 TYPE = 0 - "restore" spool (see below)
 CHAPTER = chapter number of file on source volume

Action: The request is entered in the request list on fsys FSYS, and
 if necessary an allocation request (see activity 68) is made
 in VOLUMS' name.

Reply: integer DEST, SRCE, FLAG, P2, P3, P4, P5, IDENT

 FLAG = 0 - success (N.B. means only that request entered)
 # 0 - failure

Notes:

- * USERNAME need not be the owner name found in the chapter header when the spool request is subsequently carried out. However, the file is always restored to the owner.
- * The file is created with its access permissions, except OWNP, set from the current permissions in the owner's archive index. OWNP is set from the chapter header (i.e. the OWNP in effect at the time the file was written to tape). The status bytes are otherwise not reset.
- * The requested tape must have been written in the standard format (see Section 6 below).

3. VOLUMS_services_initiated_by_operator

These services are requested by pressing the "COMMAND" key on the Operator's console and typing

V/<command>

followed by the "ENTER" key.

When the requested service does not result in a visible display, there will be an immediate response, as follows:

- DONE - The service has been successfully completed.
- OK - The command has been recognised and the requested service started. Subsequent conditions may cause the service to fail, but any such failures will be reported explicitly.
- END service - The service has ended, not necessarily successfully: there may have been error conditions which will have been explicitly reported.

The services fall roughly into three groups:

1. management of private devices and volumes (Subsection 3.1)
2. management of the backup/archive store (Subsection 3.2)
3. monitoring (Subsection 3.3)

3.1 Management_of_private_devices_and_volumes

<dev>	:= device mnemonic;	e.g. M50, M61
<devgroup>	:= <dev> ! <dev> truncated on the right	e.g. M50, M5, M
<ident>	:= volume identifier;	e.g. NRA024
<offon>	:= OFF ! ON	

Each of the commands in this group is entered on a single line.

<devgroup>?

Display, in an abbreviated form, the current status of the device(s) specified, on the main operator console; e.g.

V/M50? displays status of M50

V/M5? displays status of M50, M51, etc.

DEVS Display comprehensive status of all devices, on the auxiliary display.

<dev>=<ident>

The unlabelled volume on drive <dev> is to be known by the identifier <ident>; e.g. V/M51=NT4012

NO <ident>

The volume <ident> is not available, and outstanding requests for it are to be refused.

REQ List all outstanding volume requests on the auxiliary display.

RETAIN <dev>

RETAIN <ident>

RETAIN

The device <dev> or the device on which volume <ident> is currently mounted is to remain loaded even when the volume is not in use. If neither <dev> nor <ident> is given, all devices are "retained".

UNLOAD <dev>

UNLOAD <ident>

UNLOAD

The converse of RETAIN; this is the default.

TAPES? Display the Supervisor state of all tape drives, on the main console.

LABEL <dev> <ident>

Write the volume identifier <ident> on the volume currently mounted on device <dev>.

STIPL <dev> <ident>

Write an "S" series IPL label <ident> on the volume mounted on device <dev>.

TIPL <dev> <ident>

Write a "P" series IPL label <ident> on the volume currently mounted on device <dev>.

MODE <dev>=<offon>

Tell the Supervisor to release (OFF) or reclaim (ON) the device <dev>.

3.2 Management_of_the_backup/archive_store

<backup type> := ALL!CHANGED!WEEKLY!DAILY
(WEEKLY= all cherished files,
DAILY = all changed cherished files)

<archive type> := DESTROY!TO TAPE

<export type> := ALL!CHANGED!WEEKLY!DAILY!SELECT

<index backup type> := NEW!APPEND

<index restore type>:= ALL!null

<tape type> := BTAPE!ATAPE!ETAPE!STAPE

<fsys> := file system number (decimal)

<any or all fsys> := <fsys>|-1
(-1 means that the actual file system number is
unknown, or all file systems are to be included.
If this is the last item on the line it may be
omitted.)

<single user> := user name with all 6 characters defined

<set of users> := user name with one or more characters undefined
by '?'

<user group> := <single user>!<set of users>

<inclusive user list> := (<user group><rest of user list>)

<exclusive user list> :=)<user group><rest of user list>(

<user list> := <inclusive user list>!<exclusive user list>

<rest of user list> := ,<user group> <rest of user list>!null

<userspec> := <user group>!<user list>!null
. (null = <user group> of ??????)

<first>:=<last>:=<page>:=<pos>:= decimal number

If an inclusive user list is specified, all the users on the nominated file systems(s) who match anything in the list take part in the operation. If an exclusive user list is specified, all the users on the nominated file system(s) who do not match anything in the list take part.

Several of the commands in this group require to be entered on multiple lines; for the second and subsequent lines an appropriate prompt is output, and the response should be given without the preceding "V/".

BACKUP Initiate a backup sequence. The prompt "GIVE BACKUP SPEC" is output and a line of the form

<backup type> <any or all fsys> <userspec>

should be entered; <userspec> can be continued over multiple lines, when the prompt becomes "CONTINUE INCLUSIVE LIST" or "CONTINUE EXCLUSIVE LIST", as appropriate. Individual list items may not be split across a line boundary. The nominated files are written to tape(s) taken from the BTAPE list, a listing of tape contents is produced, the write bits for the files are cleared, the #ARCHs are updated and a secure tape is written (see Section 5).

ARCHIVE Initiate an archive sequence. The prompt "GIVE ARCHIVE SPEC" is output and a line of the form

<archive type> <any or all fsys> <userspec>

should be entered; subsequent lines are as for BACKUP. Files eligible for archive are those which have been unused for some period (see MARK UNUSED, below) or those which have been explicitly marked for archiving by the user. If <archive type> is "DESTROY", eligible files which are not cherished are destroyed and a listing is produced. If <archive type> is "TO TAPE", eligible files which are cherished are written to tapes taken from the ATAPE list, the #ARCHs are updated, a listing of tape contents is produced and the files are destroyed.

EXPORT Initiate an export sequence. The prompt "GIVE EXPORT SPEC" is output and a line of the form

<export type> <any or all fsys> <userspec>

should be entered; subsequent lines are as for BACKUP. If <export type> is SELECT a prompt of the form

SELECT user.filename?

is output for each file and each specified user, to which the response "Y" (yes), or "N" or null (no), should be given to include or exclude the file. The nominated files are written to tape(s) taken from the ETAPE list and a listing of tape contents is produced. (Note that the SELECT file confirmations are done using DOUT, and so other VOLUMS activities are suspended until the selection is complete. Since it is possible by accident or design to specify all users, a very large number of responses could be required. Two "get-outs" are available: the response "NEXT" to a file implies "N" to that file and to the remaining files for the current user. The response "END" implies "N" to that file and to the remaining files for the current user, and to all files for the remaining users.)

RESUME BACKUP
RESUME ARCHIVE
RESUME EXPORT
RESUME RECREATE

Continue a backup, archive or export sequence which was interrupted. See Section 5 for RESUME RECREATE.

RECORD BACKUP
RECORD ARCHIVE
RECORD EXPORT

If BACKUP, or ARCHIVE TO TAPE, or EXPORT, proceed to that part of the sequence (described above) immediately following the writing of tapes; if ARCHIVE DESTROY, proceed to destruction.

SUSPEND BACKUP
SUSPEND ARCHIVE
SUSPEND EXPORT

Terminate the sequence as soon as possible consistent with a resumable state; the current writing or reverse checking of a file is allowed to complete. Tape(s) which are currently being written or reverse checked have a "CHECK FAILS" forced; tape(s) which have been successfully written and checked remain valid. If load requests are outstanding, it is necessary to follow the SUSPEND by a NO <tape> in order to kick the appropriate activity.

The control files are retained and the sequence can be RESUMED. If, instead, the sequence is restarted from scratch the retained control files are discarded.

MARK UNUSED

Mark files which have not been used since the last MARK UNUSED was done. When a file accumulates a number of such unused marks it becomes eligible for archive.

INDEX BACKUP

Initiate an index backup. The prompt "GIVE INDEX BACKUP SPEC" is output and a line of the form

<index backup type> <any or all fsys> <userspec>

should be entered; subsequent lines are as for BACKUP. The non-file index information for the specified file system(s) and users is collected in a file called INDEXBACKUP (one per fsys) in VOLUMS' index. If <index backup type> is NEW, any existing file is first destroyed. If <index backup type> is APPEND there must be an existing file to which the information is appended. APPEND should be used with caution since, if a user's information appears more than once in a file, only the earliest can be retrieved by INDEX RESTORE. The INDEX BACKUP performed automatically as part of the secure operation (see Section 5) is always of type NEW.

INDEX RESTORE

Initiate an index restoration. The prompt "GIVE INDEX RESTORE SPEC" is output and a line of the form

<index restore type> <any or all fsys> <userspec>

should be entered; subsequent lines are as for BACKUP. The non-file index information contained in the file INDEXBACKUP (found in the VOLUMS index on the specified file system(s)) is used to re-accredit the specified users on that file system. If such a user is currently accredited anywhere else he is not re-accredited.

The file found on a file system need not have originated there. Thus users can be re-accredited on a file system other than their original.

RELOAD <ident> <first> <last>

Spool the files in chapters <first> to <last> inclusive on tape <ident> into the on-line file system. If neither <first> nor <last> is given, all chapters on the tape are implied; if <last> is not given, the single chapter <first> is implied.

The files are created with access permissions, arch byte, codes byte, cct and ssbyte set from the tape file header (i.e. as they were at the time the tape was written). Except for VOLUMS and SPOOLR files, which are re-created on the file system where they originated, all files are created on the file system where the owner is currently accredited.

TRANSFER <ident> <first> <last>

As RELOAD (above), except that existing files of the same names in the destination file system are first destroyed, and the status bytes are not reset.

REPLACE <ident> <first> <last>

As RELOAD (above), except that there is a prompt for the destination file system. The value given is used instead of any which might be in the tape file header.

If a file's owner is not currently accredited on the specified file system the file is not spooled.

RECREATE <fsys>

Initiate a re-creation with <fsys> as the destination file system. The prompt "GIVE TAPE SRCE USERSPEC" is output and a line of the form

<ident> <fsys> <userspec>

should be entered; subsequent lines are as for BACKUP. <ident> is the identifier of the secure tape to be used and <fsys> is the source file system on that tape. See Section 5 for a detailed description of RECREATE.

DISCARD BACKUPS

For all file systems and all users, discard #ARCH backup entries which have expired. See Section 5 for a detailed description.

REQLIST <page>

List page <page> of outstanding spool requests, on the auxiliary display.

BADREQ <ident> <first> <last>

Delete, from the request lists, the spool requests for tape <ident>, chapters <first> to <last> inclusive.

NEW <tape type>LIST

Create a new tape list of the specified type on the current IPL file system (see Section 4). The list is created empty.

ADD <tape type> <ident> <pos>

Add the identifier <ident> to the specified tape list at position <pos>, where $1 \leq \text{<pos>} \leq \text{current length of list} + 1$.

If <pos> is not specified the end of the list is assumed.

REMOVE <tape type> <ident>

Remove the identifier <ident> from the specified tape list.

<tape type>S

Display the specified tape list on the auxiliary display.

3.3 Monitoring

Each of the commands in this group is entered on a single line.

PRINT Print the current VOLUMS log on the line printer and initialise a new log file.

MONON Switch on monitoring.

MONOFF Switch off monitoring.

STOP Stop.

4. Management of the BACKUP, ARCHIVE, EXPORT and SECURE tape lists

<tape type> := BTAPE ! ATAPE ! ETAPE ! STAPE

The four tape lists reside in files VOLUMS.<tape type>LIST. They are created and initialised on the current IPL file system (or, if the file already exists on that file system, reinitialised) by the NEW command; see Subsection 3.2.

Tape identifiers are added to a list by ADD and removed from it by REMOVE.

Note that when adding a tape identifier to a list, or retrieving one for use, the available file systems are searched from 0 to find a file with the appropriate name. It is thus important that the first one found is the current one, and in practice there should only be one list file of each type across all file systems. If a new list file is created on a file system other than that which contains the current one, the latter should be subsequently deleted by MANAGR. (VOLUMS does not attempt this since the file system may be off-line.)

The current contents of a tape list are displayed on the auxiliary display by the <tape type>S command.

The identifiers of tapes to be written are taken from the front of the lists. In the case of backup or secure tapes they are immediately replaced at the back; i.e. these lists are cyclic. These two lists also contain, for each tape, the date when it was last written (see Section 5).

The ETAPE list used by EXPORT is not managed, and the ETAPES and ADD commands should be used prior to the export to ensure that the required tapes are at the front of the list. It would be good practice to leave this list as it was found, if necessary using the REMOVE command.

5. File System Re-creation

This section includes descriptions of archive indices, secure tapes, the backup tape cycle and backup tape list (BTAPE list), and index backup and restoration insofar as they contribute to file system re-creation.

Each user has an archive index (#ARCH) which contains, in addition to archive entries, entries for all the user's files on backup. Backup tapes are used cyclically, and so these backup entries have only a temporary validity. Since a #ARCH containing references to a particular backup tape may be unavailable (e.g. off-line) when that tape is about to be rewritten, it is not possible (nor would it be efficient) to search for and discard such backup entries at that time. Instead, a backup entry is validated by comparing its date with the date when the tape was last mounted with a write ring (this is held in the BTAPE list). If they are not equal the entry is no longer valid. Discarding such expired entries is thus a housekeeping operation which can be performed at any time using the command DISCARD BACKUPS.

Thus, given

- 1) the BTAPE list,
- 2) the set of users' #ARCHs, and
- 3) the backup tape cycle

it is possible to re-create automatically all users or a group of users, on one or more file systems. This has been implemented as follows.

Following the daily backup, when all available #ARCHs have been updated, an INDEX BACKUP(NEW) is automatically performed which collects users' non-file index information in a file called INDEXBACKUP (one per fsys) in VOLUMS' index. Then a further tape, called a "secure tape", is written containing the BTAPE list, and for each file system this index backup file followed by the users' #ARCHs in the same order as the users' entries occur in the index backup file.

To re-create all users or a group of users on a destination file system, it is only necessary to nominate the secure tape to be used and the source file system on that tape. Thus:

V/RECREATE <destination fsys>

after which the prompt "GIVE TAPE SRCE USERSPEC" is given.

A search is first made for the BTAPE list on all on-line file systems. If it is not located it will be retrieved from the secure tape during the following procedure. A scan is made of the secure tape to locate the index backup file for the source file system. This file is REPLACed on the destination file system and used to do an INDEX RESTORE for the specified users, which re-accredits them and puts back their non-file index information (password, delivery information, etc). Then the #ARCHs for the same users are located on the tape (immediately following the position of the index backup file) and REPLACed on the destination file system. Any missing #ARCHs, unavailable at the time the tape was written, are detected and reported at this time.

When all the available #ARCHs have been REPLACed, the backup entries in them for the specified users are processed to obtain the entries for the most recent versions of all their files on backup, and ordered to be

retrieved in a single pass through each required backup tape. Standard REPLACE requests for all the necessary tape chapters are then generated. (Only a single request is generated for adjacent chapters, but there could still be a very large number of requests. The request lists on all on-line file systems are used if necessary, but if all these fill up the generation is suspended and can be resumed when space becomes available by the command RESUME RECREATE.)

Since only the most recent versions of files are requested, the order of mounting the backup tapes is not important. They are requested in ascending order of the tape sequence number, and the REPLACES may proceed in parallel, using as many tape drives as are available.

The RECREATE operation is officially complete when all the REPLACE requests have been generated, since these are standard requests and will be preserved across a crash. The system could be opened to the users at this point since they are accredited, their files gradually returning as the REPLACE requests are actioned.

If more than one file system is to be re-created, the next one can be started at this point. The REPLACE requests generated for it will join those already in the request lists and the normal mechanism for selecting the next request to be actioned will effectively merge them.

Note that the scan of the secure tape is triggered by a single REPLACE request for all chapters on the tape, and the integrity of the scan depends on there being no other requests (e.g. RELOAD, TRANSFER, etc.) for the secure tape while the scan is in progress.

Secure tapes to be written are taken from a small cyclic list called the STAPE list. A listing of the secure tape contents is produced as for backup etc., once it has been successfully written and checked. Writing the secure tape is considered to be part of recording the backup. Thus there is no explicit command to initiate it. If, following a successful backup, the secure tape is not written for some reason, this can be achieved by RECORD BACKUP. (This will note in passing that the BACKUP control files do not exist.)

A full scale RECREATE consists of a sequence of operations which can be initiated in isolation if required, as described below. Notice, however, that it is always safe to do a complete RECREATE, since none of its operations are destructive: if a user is currently accredited anywhere, he will not be re-accredited on the destination file system; existing on-line files will not be overwritten by any coming from backup, since REPLACE is non-destructive. Even if the source or destination file system is specified incorrectly, the worst that could happen is that user-destroyed files might reappear on the destination file system.

If the secure tape identifier is given as "SKIP", RECREATE assumes that the required users are already accredited on the destination file system and have up-to-date #ARCHs. It thus proceeds directly to the processing of backup entries from the latter, bypassing the scan of the secure tape and the index restoration.

Index backup and restoration can be initiated by the commands INDEX BACKUP and INDEX RESTORE. No tape transfers are involved: the non-file index information is gathered into or dispersed from the on-line index backup file in VOLUMS' index.

Such partial re-create operations may be useful in a variety of situations; e.g. when a file system is only partially corrupt, or when #ARCHs are missing from the most recent secure tape, or when the secure tape cannot be read successfully (in which case a semi-manual re-create must be done, perhaps involving the next most recent secure tape.)

EHAS_2900_Standard_Tape_Format

Tape format: 80 byte standard VOL1 label
TM
Chapter 1 (blocksize = epage size)
TM
Chapter 2
TM
.
.
.
Chapter n
TM
TM

Chapter format: Block 1 header page (see below)
Block 2
.
.
.
Block n+1 } n epages of file

Header page format:

string(6) TAPENAME, USERNAME, string(15) FILENAME, string(8) c
DATE, TIME, TYPE, byteinteger SPARE0, SPARE1, SPARE2, integer c
CHAPTER, EPAGES, FSYS, PERMS, OWN, EEP, ARCH, CODES, SSBYTE, c
CCT, SPARE3, SPARE4, SPARE5, RECORDS, string(6) OFFERED TO, c
recordarray PERMLIST(1:16) (string(6) USER, byteinteger PERM)

TAPENAME	name of the tape this file is on
USERNAME	owner of the file
FILENAME	name of the file
DATE	date when written to tape
TIME	time when written to tape
TYPE	"BACKUP", "ARCHIVE" or "EXPORT", as appropriate
CHAPTER	the chapter number on the tape
EPAGES	the number of epages in the file, not counting the header
FSYS	-1 for all but "VOLUMS" and "SPOOLR", when it is the original file system of the file
PERMS	1 if access permissions are present, else 0
OWNP	owner's access permission at time of writing

EEP	everyone else's permission at time of writing
ARCH	arch byte
CODES	codes byte
SSBYTE	ssbyte
CCT	connect count
RECORDS	number of individual permissions (in permlist)
OFFERED TO	the user to whom the file is on offer, else null
PERMLIST	array of individual permissions

C.D. McArthur