Pre-history: the origins of HSM in scientific computing in Australia

Robert C. Bell | CSIRO IMT Scientific Computing

29 August 2019 Updated to 12 July 2019

Outline

- Terminology
- History – CSIRO – the start of on-line storage
- History – CSIRO – succession of services
- History – issues for today

- (Similar stories for other organisations)

Terminology

- **Transfer of data between primary and secondary media**

Process	**'Disc'**	**'Tape'**
Backup	data and metadata remain	copy of data and metadata made
Archive	data and metadata removed	copy of data and metadata made
Migration (HSM)	data removed, metadata remains	copy of data made (and sometimes metadata)

From Barbara Ainsworth:
Pearcey and CSIRAC
Pearcey's centenary this year

- Pearcey proposed a project in 1946 for CSIR to develop a digital general purpose computer.
- He worked with Maston Beard and their team to build CSIR Mk1 (later called CSIRAC); the first successful sequence of operations under program control achieved in November 1949.

CSIRO ScienceImage no.3539 CSIR Mk1 in 1951

Dinosaurs from 1949

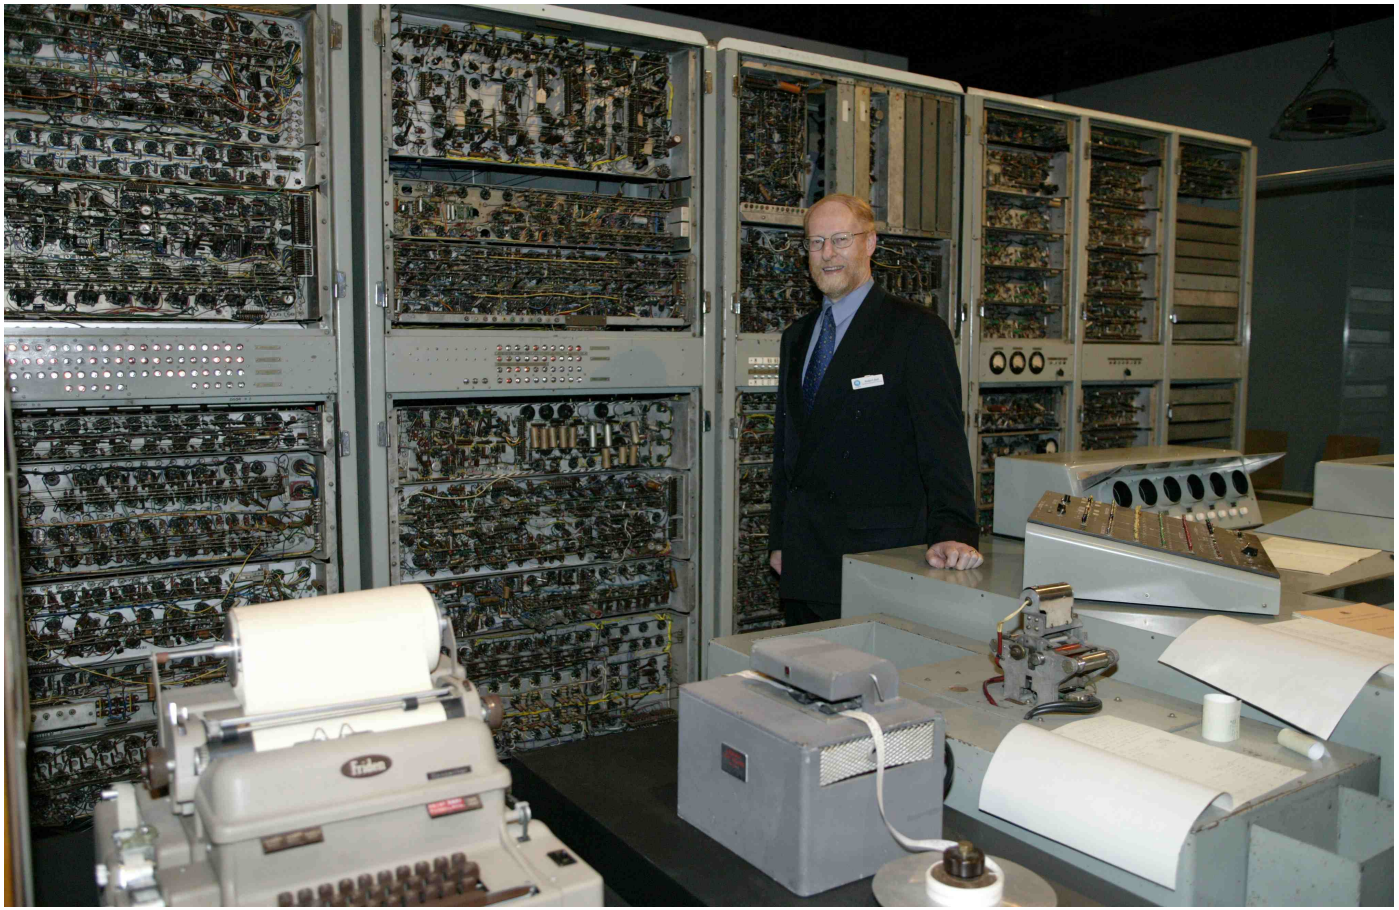

From Barbara Ainsworth: Pearcey and CSIRONET

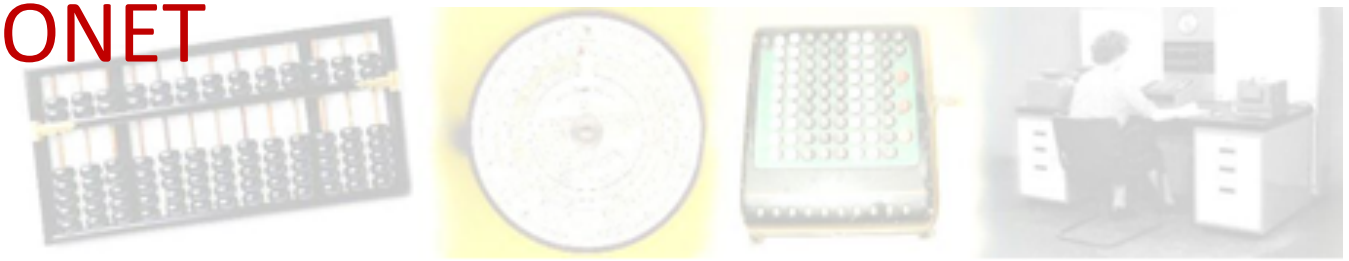

- Pearcey proposed that the CSIRO should have a network of computers of varying levels of computer capability with a set of CDC computers at CSIRO research laboratories.
- Project commenced 1963 with a CDC 3600 in Canberra and CDC 3200s in Sydney, Melbourne and Adelaide.
- All the machines had a FORTRAN facility to ensure the uniformity of programming language by users. Each local installation could receive tasks expressed in a common language. Information transferred physically and then electronically from 1966.

CSIRO Film - Classifying by Computer 1966

History: CSIRO's Computing Services

- In 1962, a proposal from Trevor Pearcey with the support of E.A. Cornish from DMS for the establishment of a 'network' of computing facilities was approved by the CSIRO Executive and a £3 million grant (about $88M in 2019 values) was provided by the Commonwealth for the establishment of a Computing Research Section within the CSIRO (Pearcey 1988).
- CSIRO Computing Research Section – 1 January 1963 in Canberra
- Tenders for supply of equipment:
  - CDC 3600 in Canberra, satellite CDC 3200 systems in Sydney, Melbourne, Adelaide
  - Remote batch processing – cards and printouts – air freighted tapes
  - "The Scientific Computing Network"

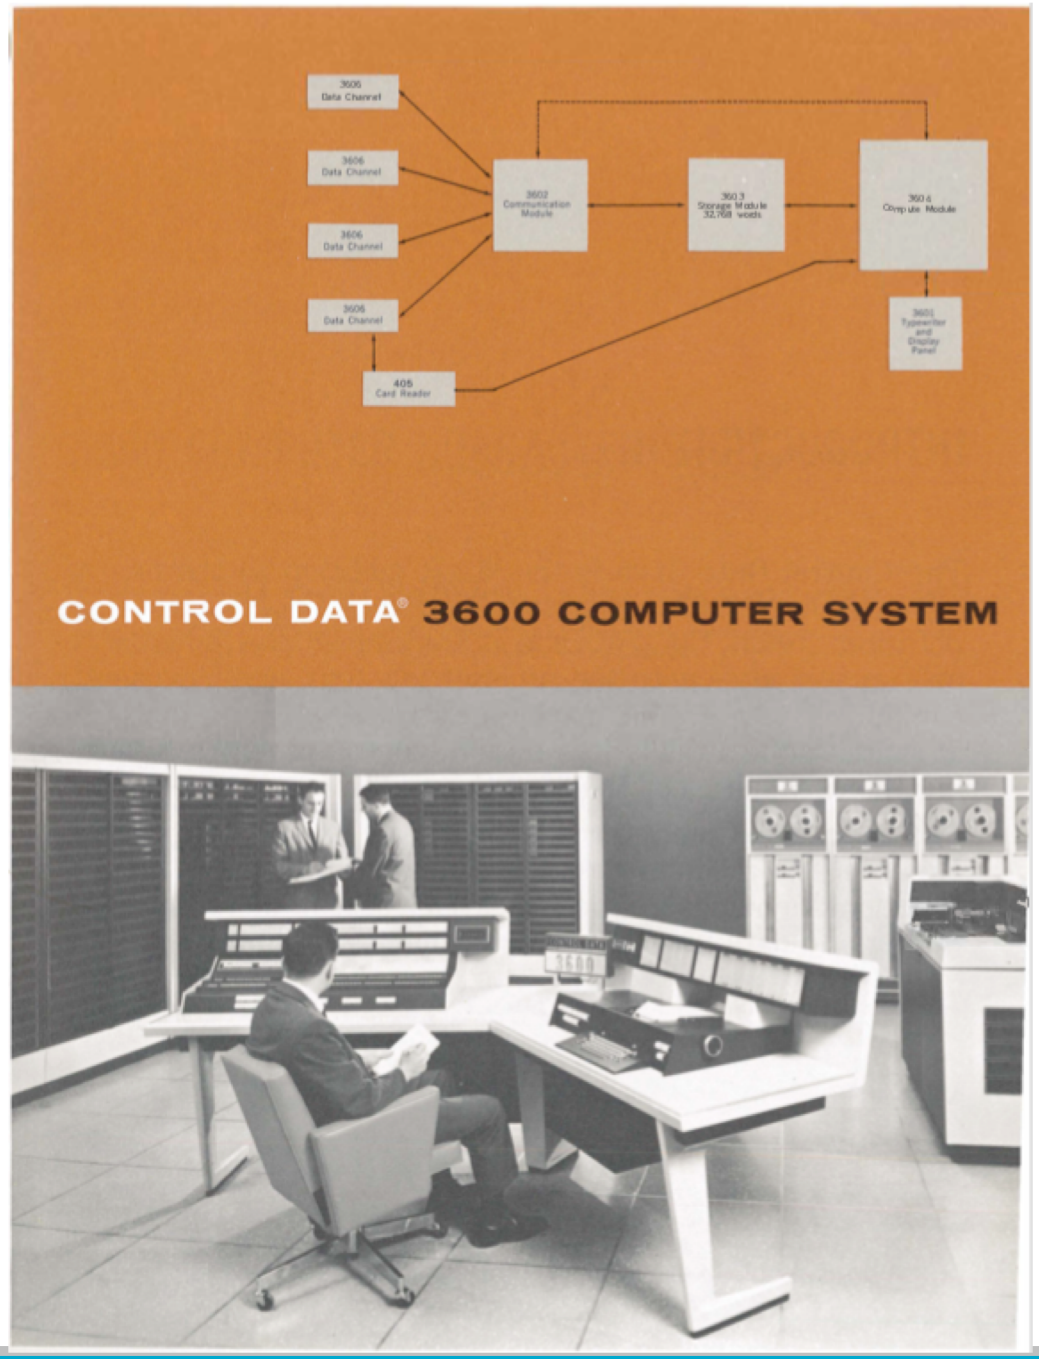
3606
Data Channel
3606
Data Channel
3606
Data Channel
3606
Data Channel
405
Card Reader
3602
Communication
Module
3603
Storage Module
32,768 words
3604
Compute Module
3601
Typewriter
and
Display
Panel
CONTROL DATA® 3600 COMPUTER SYSTEM

History – on-line storage

- Focus on computing power: initial systems had cards or paper tape for input, line printers for output, and intermediate storage on magnetic tape – even compilers, loaders, subroutine libraries were on tape only.
- Intent was to acquire a disc, but a drum unit was instead acquired first and arrived in March 1966.
- 4 units, each holding 500,000 48-bit words, for a total of 12 Mbyte
- Transfer speed of 1.5 Mbyte/s (not clear whether that is per unit or total). So time to read entire contents was either 2 seconds or 8 seconds (c.f. half a day to read the entire contents of a modern disc).
- Needed new OS
  - Drums And Display (DAD) written by CSIRO and CDA staff.

Example Magnetic Drum

http://forpresentation12.blogspot.com/2015/01/magnetic-drum.html

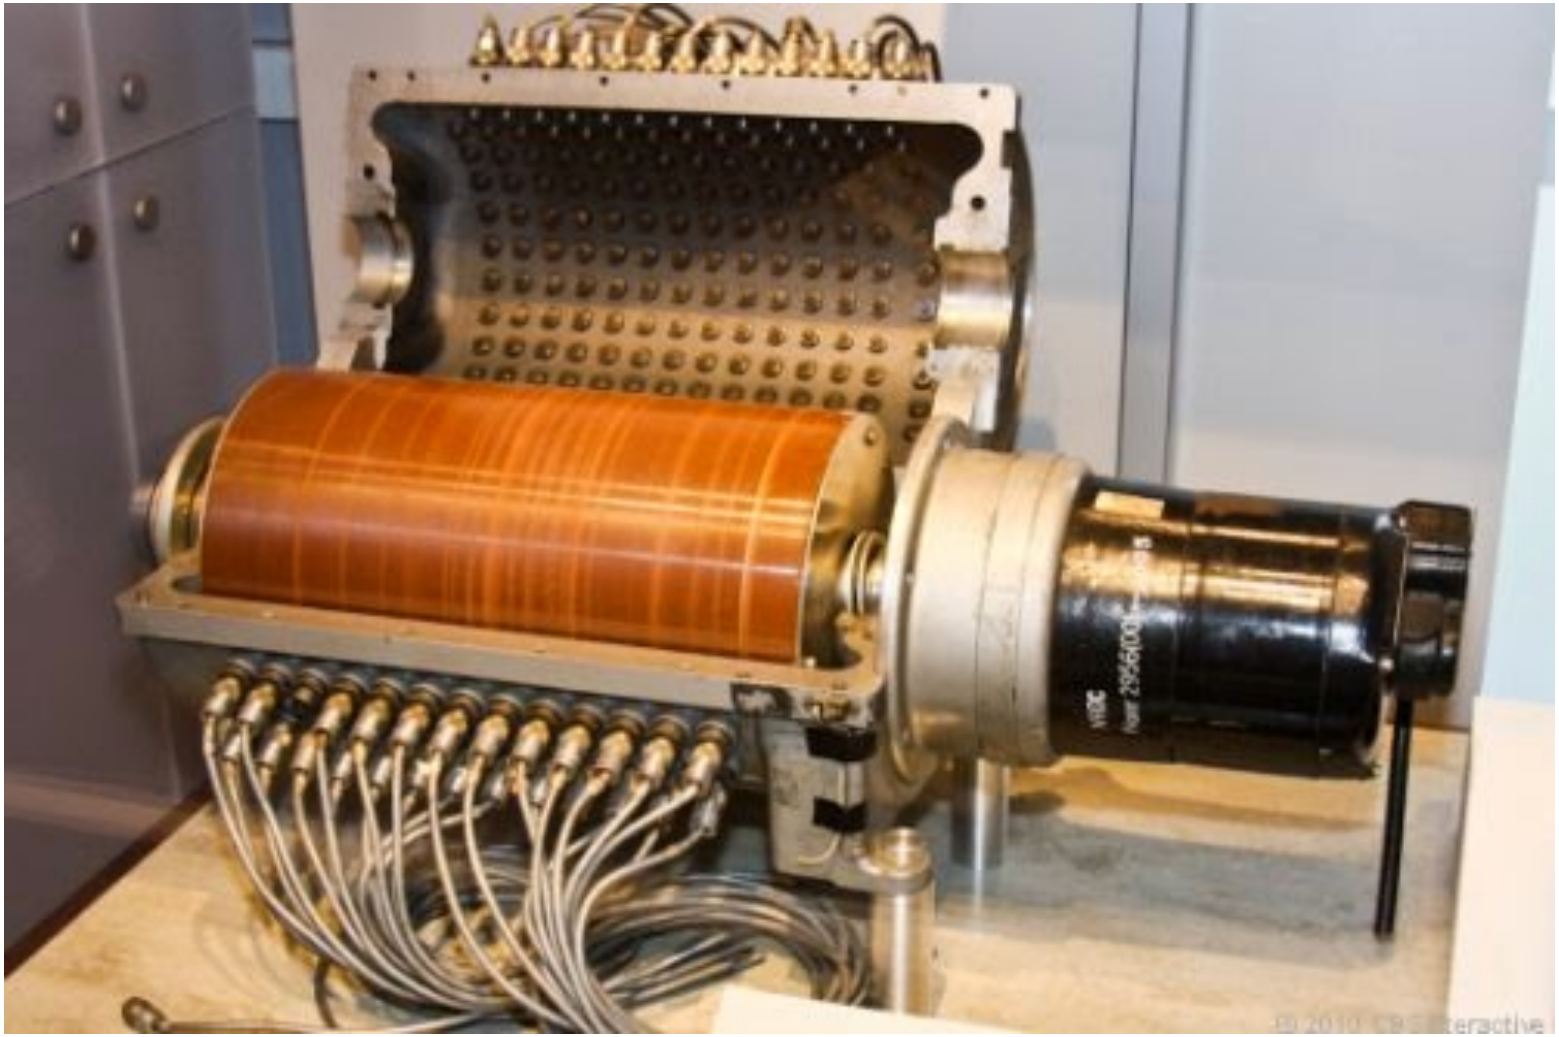

History – on-line storage

- DAD – Document Region
- First on-line storage for CSIRO
- Used primarily for systems and libraries, small amount for users.

DAD

The basic item of information is a document which can be for example a program, a sub-progra, data, program and data, etc. In general a stack of documents will be read by the card reader and whatever their final requirements the information from these documents will be written on to the magnetic drum under the control of a background program. Similarly input from the paper tape or magnetic tape stations and from the keyboard consoles can be routed to the drums by the background program. These documents if required for execution as indicated on the relevant control cards will be entered on to one of two Execution Lists either EL1 or EL2. Short running jobs i.e. less than K mins., (K at present is 5) will be placed on EL1 which will have priority over all the jobs in EL2.

DAD

When the DAD system is fully operational the user will be able to retain documents containing data or relocatable subroutines or Fortran subroutines on the drum. Using the keyboard consoles he will be able to edit old documents, create new documents and call documents for execution.

DAD

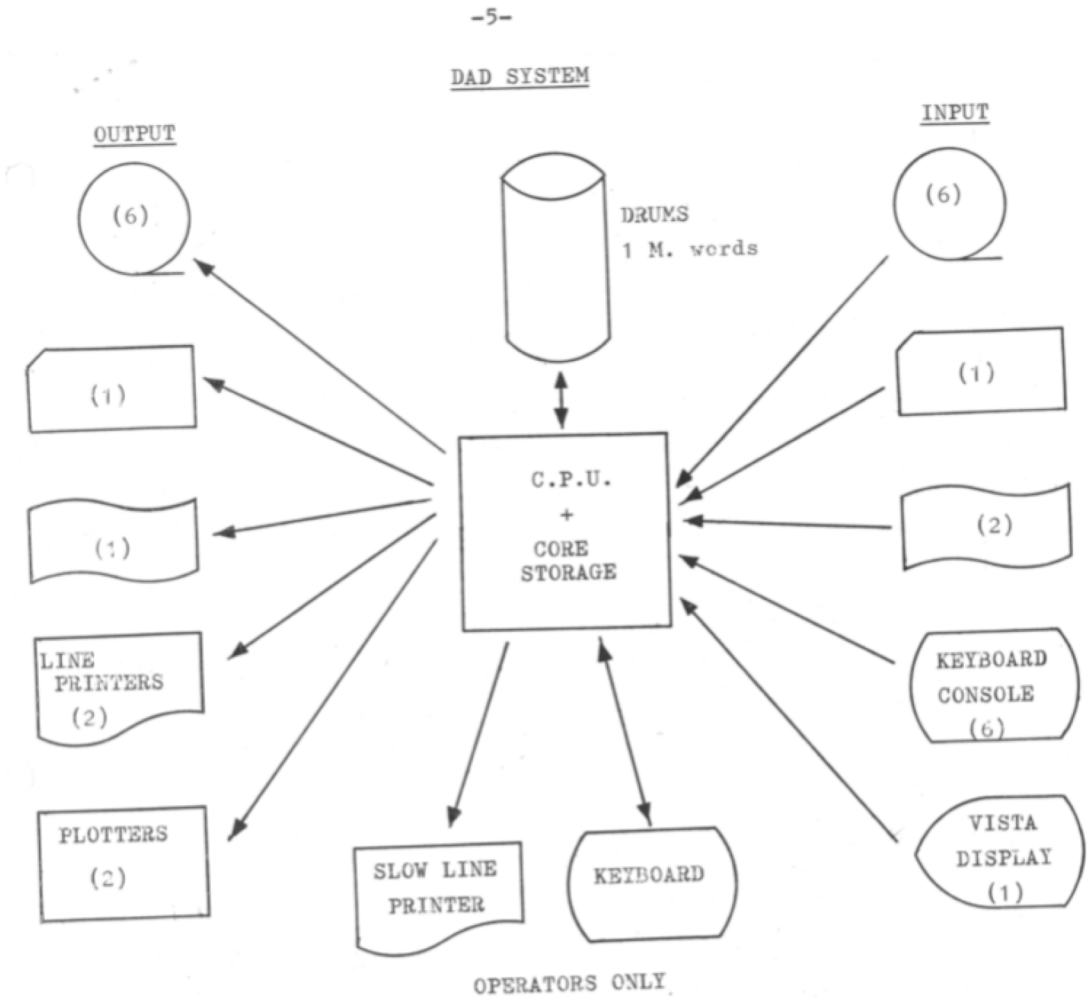
-5-
DAD SYSTEM
OUTPUT
INPUT
(6)
DRUMS
1 M. words
(6)
(1)
(1)
C.P.U.
+
CORE
STORAGE
(1)
(2)
LINE
PRINTERS
(2)
KEYBOARD
CONSOLE
(6)
PLOTTERS
(2)
SLOW LINE
PRINTER
KEYBOARD
VISTA
DISPLAY
(1)
OPERATORS ONLY

DAD

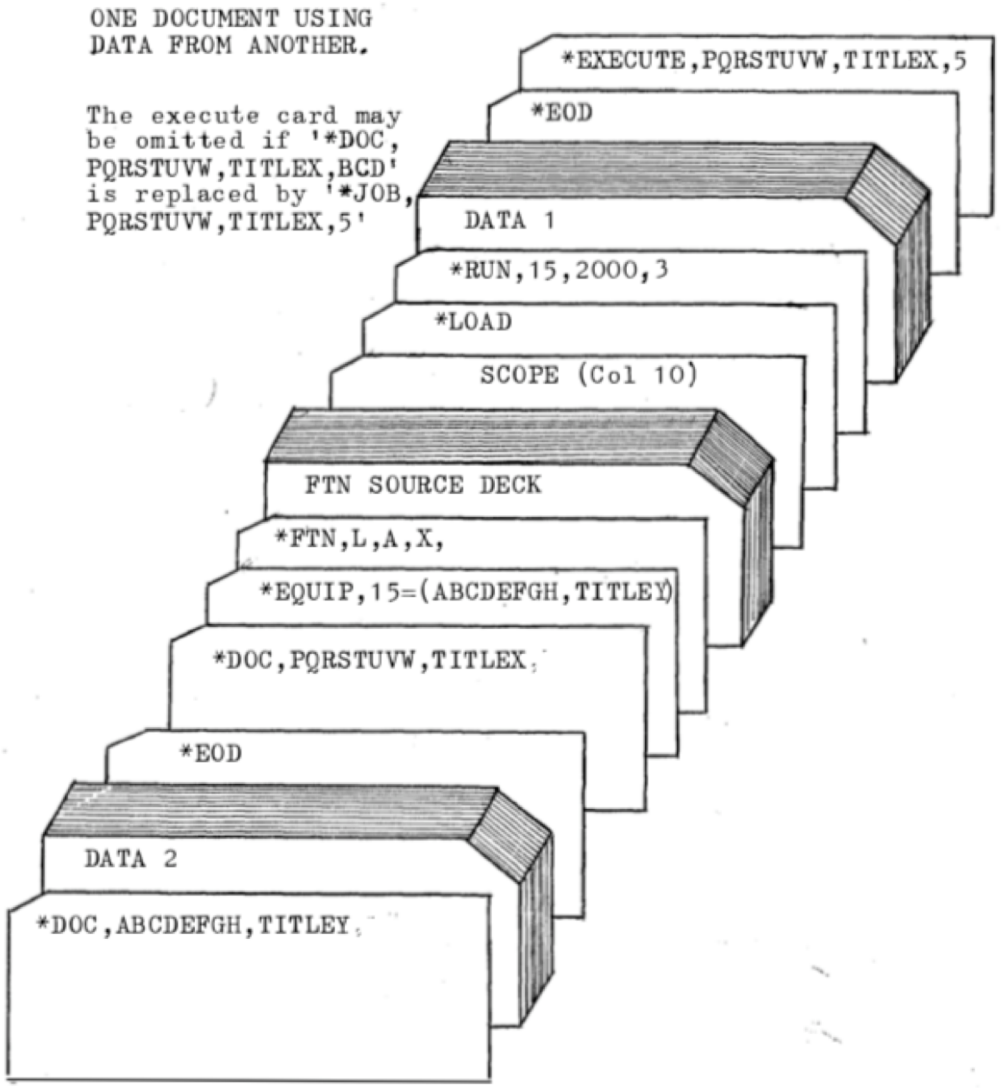

N.B. *EXECUTE statement not implemented in the initial version of DAD for CR input. It is available via the Keyboard displays.

DAD
March
1967

Saving Documents Future Policy

It is expected that within a few months the system will be able to save documents for extended periods. It will then become necessary, at times, to discard documents when the drums are full. Documents which have names with the last two characters in the identification field numeric and in the range 01 to 69 will be discarded first.

Charging: set in advance: $108 per hour for compute, $7.50/Mbyte/day.

Disc: October 1967

Disc Files

Computers in the C.S.I.R.O. network are presently having their total storage capacity increased by the addition of disc files. Some notes on the use which it is planned to make of these facilities is given below for both the 3200 and 3600 systems.

3600 Disc System

The 3600 computer in Canberra is being equipped with a disc controller and Type 813 non-interchangeable disc file of 100 million characters capacity (i.e. 12½ million 3600 words).

It is proposed to make approximately 10% of this available to DAD users as a random access bulk store, similar to the random access facility for the drum. The remainder of the disc file will be used by the DAD system as a bulk store of serial documents. It is a design aim that documents on the disc should survive most DAD system failures, and that, in any case, it would be most unlikely that all documents would be lost as a result of a single failure.

Managing on-line storage for users

- Drum was unreliable – crashes and preventative maintenance
- Flushing arises! A continuing problem.

March 1968

Display Program DISCDOCS

DISCDOCS has been added to the 3600 display program library as a general purpose means of allowing users to retain documents from day to day in the disc file system. At present the facility is only available to display users, however, software is being developed to extend the facility by means of control card options in main jobs. As the system is still under development, no guarantee is given at the present time that documents will remain on the disc, and currently they are not saved over the weekend.

August 1968

II. 3600.

Disc Saving

Formerly, disc documents were retained during the week only, maintenance at weekends causing loss of these documents.

Now the disc documents are transferred to magnetic tape storage at the beginning of this weekend maintenance period, and are restored at its end according to the following algorithm:

At disc shutdown, all disc documents are sorted on the date of most recent activity, and saved on magnetic tape.

During disc startup, documents are returned to the disc in order of most recent activity, until all documents are restored, or until the disc reaches a pre-determined level of capacity.

Feb 1969 – flushing

II. 3600

Due to the increasing use of the disc, and the attempted indefinite retention by users of disc documents, it has been found necessary at times to expunge unaccessed documents of earliest retention date.

Currently at the close of the week's processing all disc documents are sorted into order of latest reference date, and these documents are copied to a number of magnetic tapes. A full disc requires about 7 magnetic tapes.

At the start of the next week's processing four of these tapes are copied back to the disc. The remainder of the tapes are held for a short period and then reused.

The retention time of disc documents remains undefined, as it depends on total disc activity. Recently it has ranged from 6 to 10 weeks.

It is hoped that in the near future it will be possible to transfer documents between the drums and magnetic tapes. In the meantime it is advisable for the user to write his document on magnetic tape if he wishes to retain it for an indefinite period but does not intend to make frequent use of it.

Flushing

- Dealing with filling filesystems
- Files sorted by date of last access, and older files discarded.
- Retention period not defined – depends on activity – right solution, rather fixed age currently still practised.
- July 1969

Document Flushing

All documents held on the magnetic drum have an entry in the Main Document List to identify them. One word of this entry has contained the date of last activity. The format of this word has been changed to include the time of last activity also. At present FLUSH makes use of this new information to avoid deleting any document less than 15 minutes old. This use of the time by FLUSH is only tentative at present and will probably be improved.

Feb 1970

(a) Rental for storage of disc documents

The term 'rental of storage' used in the circular refers only to inactive disc documents, that is, documents which have not been used within a period of one day. The mechanism to be adopted is that all disc documents will be scanned once a day, and those with an activity date and time more than 24 hours earlier than the scanning time will be charged one day's rental. Where the computer is not being operated, for example over a weekend, the daily rental will not be charged.

It can be seen that 'shared files', that is, documents placed on the disc by one user for his own benefit and that of others, will not have rental charged to the originator if other users are making use of such documents.

It should be noted that from the date of application of the new charges the command

*DFLOCATE, parameters

will not update the activity date.

Feb 1970

- Files in active use were not charged.
- The ‘touch’ command was disabled!

Sep 1970

AUTOMATIC DISC FLUSH 3600

As from 21 September, the retention of documents on the disc will again be subject to automatic flushing procedures. The flushing procedures were suspended after charges introduced for storing inactive documents on the disc in February this year produced a sharp drop in the use of the disc for long term storage. Since then no documents have been deleted. However, the document storage area on the disc is now becoming full, and it is necessary to clear space for new documents.

At the beginning of each week, all documents of age greater than two months will be flushed from the disc. (The age of a document is defined as elapsed time since it was last copied from the disc to the drum or since it was created, whichever is the lesser.) Further documents will then be deleted, starting with the oldest, until sufficient space remains for new documents to be added during the week. Initially it is expected that it will not be necessary to delete documents less than two months old, but as disc usage increases the life expectancy of an unreferenced document will decrease.

Users are reminded that a 'card' document on the disc may be punched out on cards by submitting the following job.

```
*JOB,charge code,ident,time
*ERASE,62
*DFCOPDR,62,,docname
*EOD
```

Documents that cannot be punched on cards may be saved on magnetic tape by using UTILITY (see INSTANT DAD).

Over-flushing done!
Two month maximum age – then flushing to reach threshold
Can still make card decks!
Later – two days grace period – allowing for delayed airline flights

March 1971 - Turnaround

As a result, turnaround of larger jobs suffers some delays. Some guidelines as to turnarounds for various estimated job times, under current load conditions, can be given:

Job time	Turnaround
breakin	Immediate - Some minutes
1-2 minutes	≈1 hour
3 minutes	< 4 hours
4 minutes	Several hours
5-15 minutes	overnight
15-30 minutes	2 days
30-60 minutes	2-4 days
> 60 minutes	weekend

March 1971

A project in hand will shortly be introduced as a service to save flushed documents on magnetic tape and to render them accessible to the user for a period of eight weeks after flushing. References made to flushed documents will then result in a diagnostic

DOCUMENT SAVED ON TAPE DATE. . . . . .

or similar, to the user. Details of this service will be announced shortly.

The beginnings of an HSM!

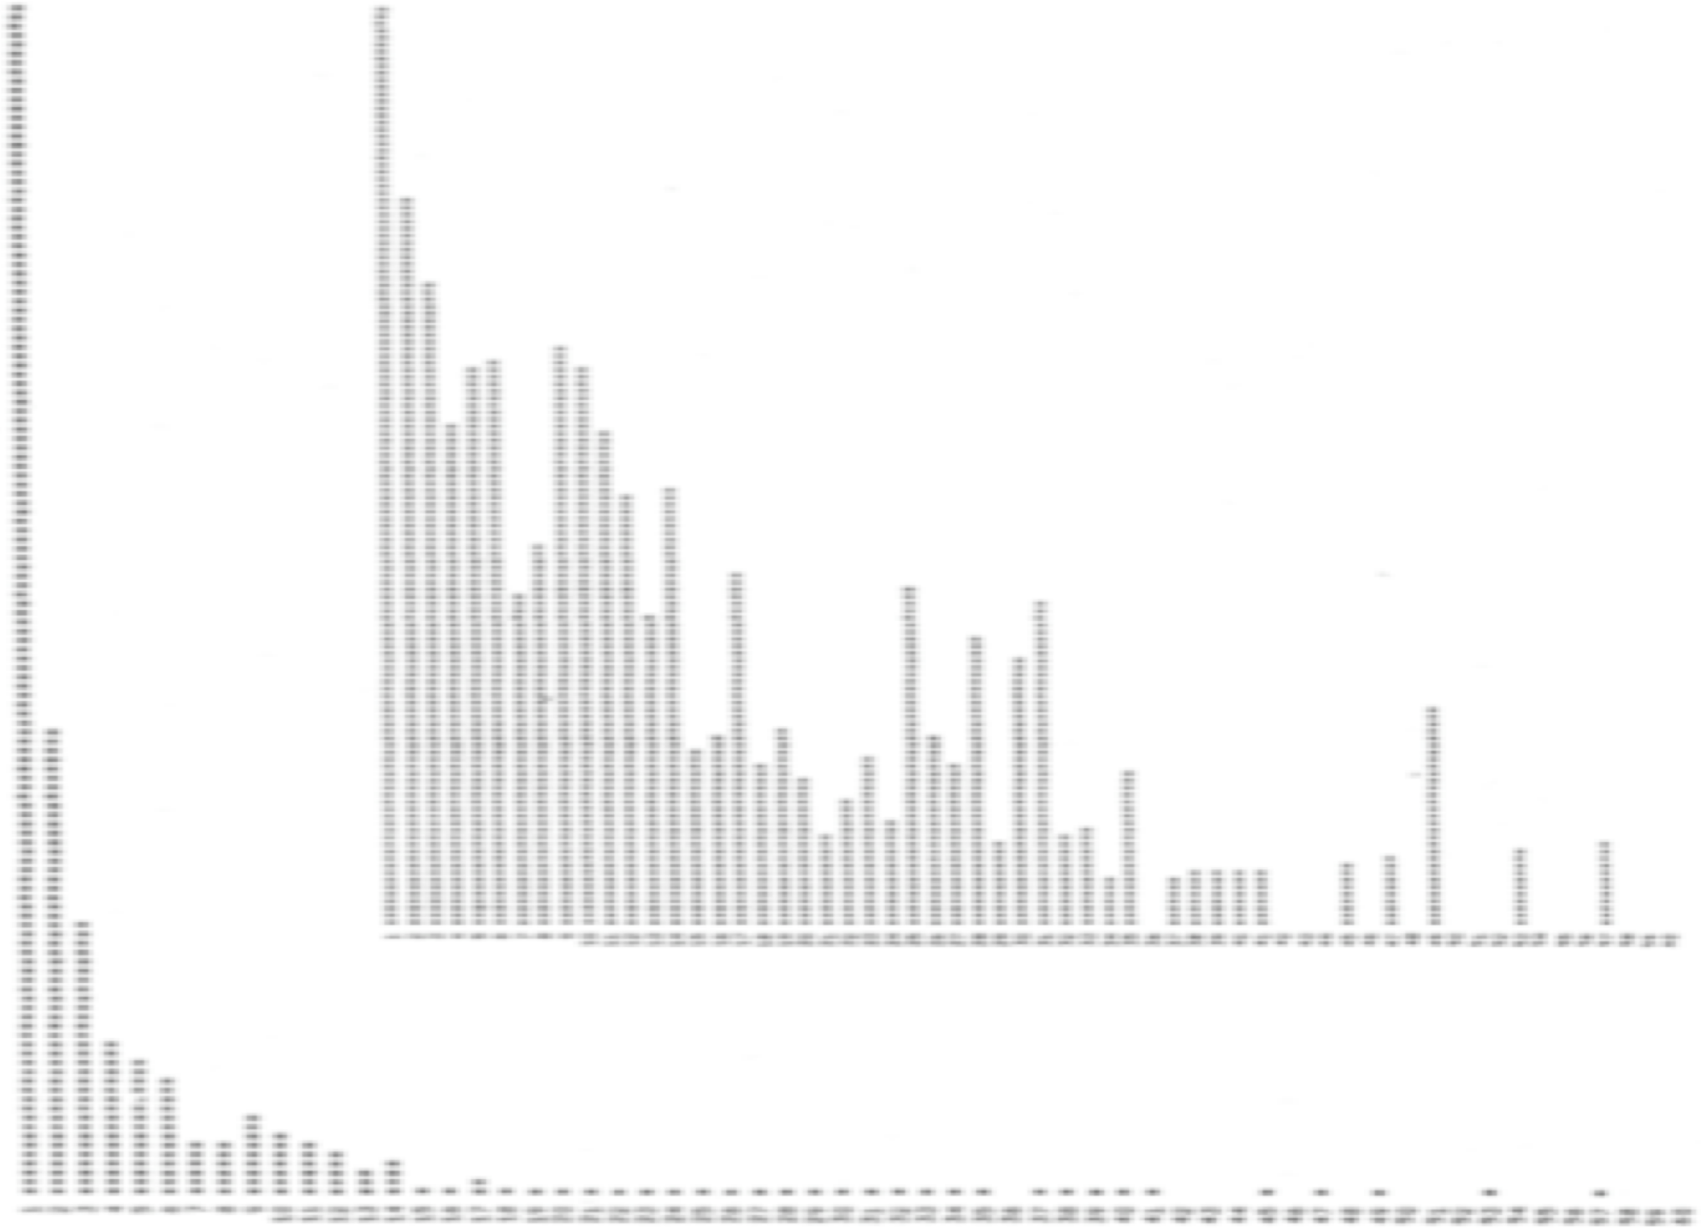

May 1971

DISC DOCUMENT FLUSHING 3600

From the beginning of this month documents flushed from the disc will be held on magnetic tape for a period of eight weeks. It was decided to introduce this back-up service following the abrupt change that took place two months ago in the ages of unreferenced documents left on the disc after weekly flushing procedures (see Newsletter 67).

A set of DFSAVE tapes will be maintained in the machine room, each tape having a label of the form

DFSAVE*dd/mm/yy*

where *dd/mm/yy* is the creation date of the tape.

An attempt to access a flushed document using the disc control statements DFCOPDR or DFLOCATE will result in the diagnostic

DOCUMENT SAVED ON TAPE LABELLED DFSAVE*dd/mm/yy*

and processing continues. (A call to DISCDOCS from a Fortran program will return the value 11.)

May 1971

- Three levels of persistent storage:
  - Drum, disc, tape
- Commands to audit files and to move/copy files between the levels.

- Feb 1972
- Unified Document Storage
- ”automatic archival and retrieval of documents to magnetic tape will be implemented”

Feb 1972

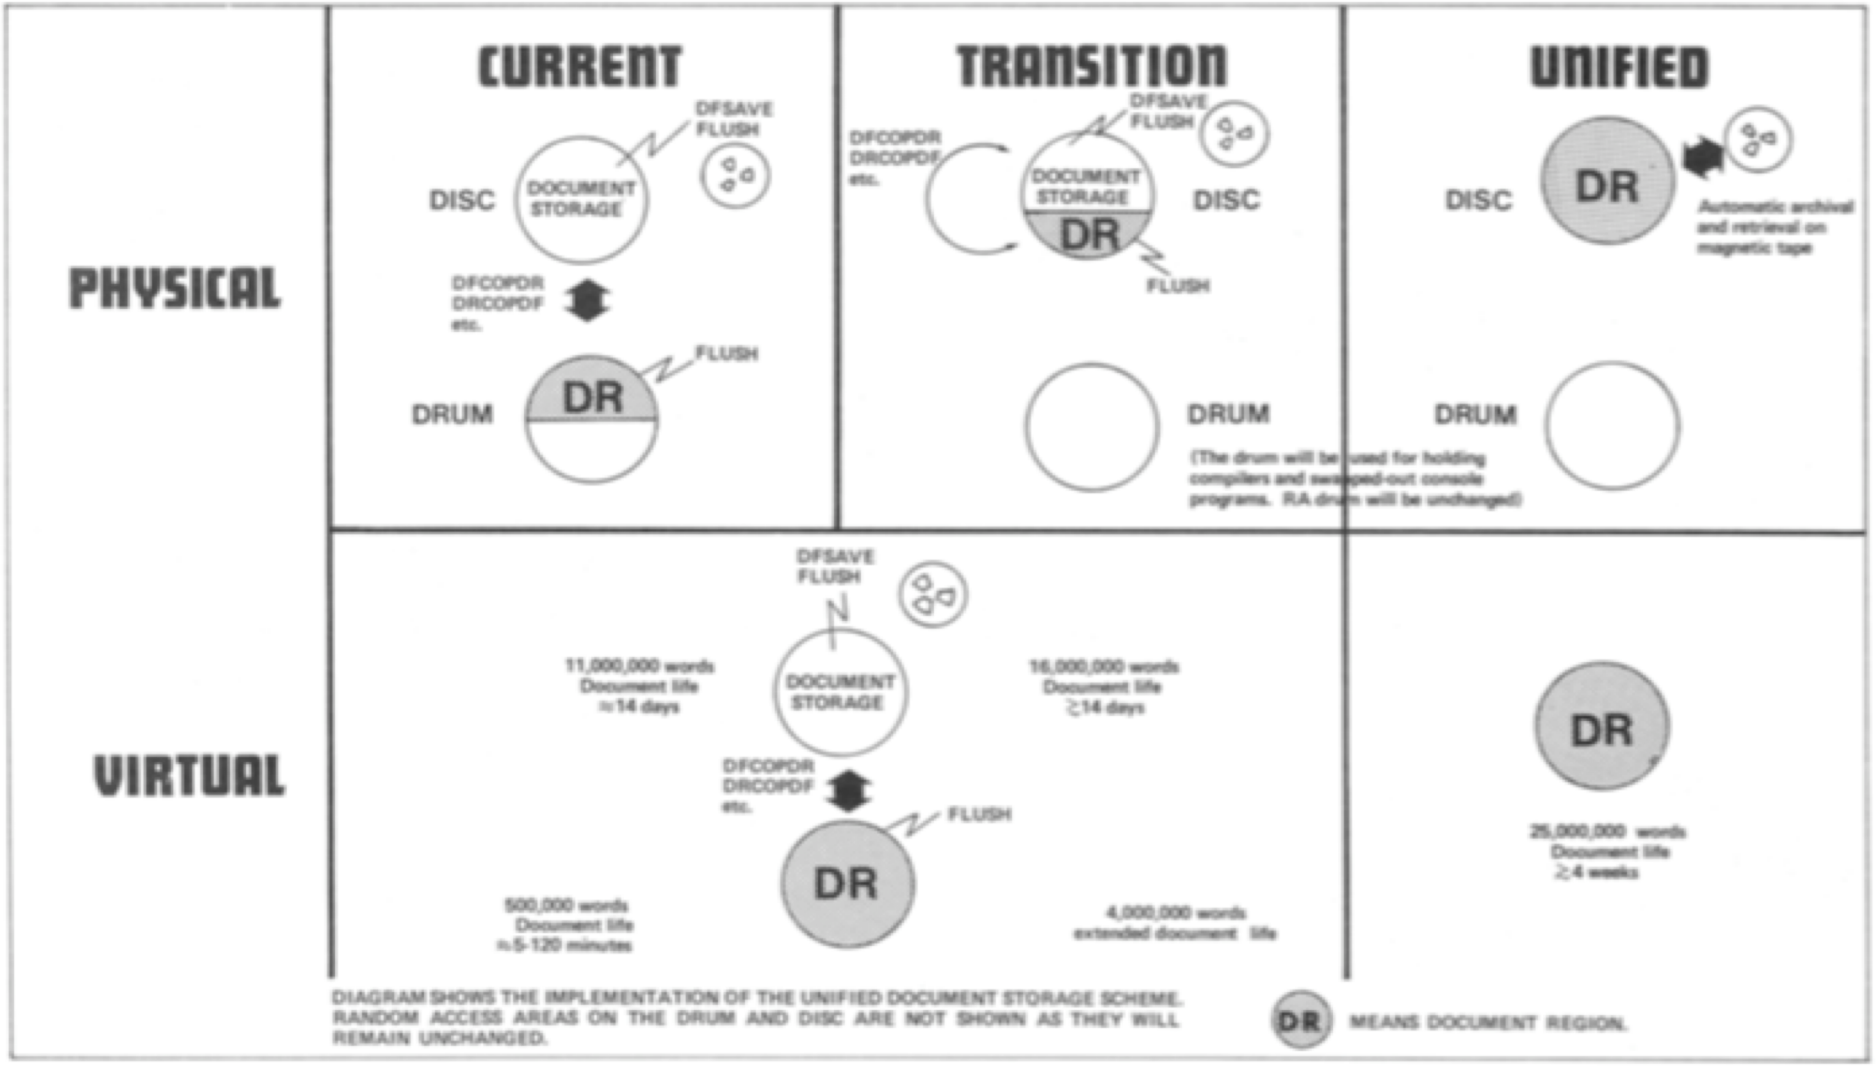
CURRENT
TRANSITION
UNIFIED
PHYSICAL
VIRTUAL
DISC
DOCUMENT STORAGE
DFSAVE
FLUSH
DFCOPDR
DRCOPDF
etc.
DRUM
DR
FLUSH
Automatic archival and retrieval on magnetic tape
(The drum will be used for holding compilers and swapped-out console programs. RA drum will be unchanged)
11,000,000 words
Document life
≈14 days
16,000,000 words
Document life
≥14 days
500,000 words
Document life
≈5-120 minutes
4,000,000 words
extended document life
25,000,000 words
Document life
≥4 weeks
DIAGRAM SHOWS THE IMPLEMENTATION OF THE UNIFIED DOCUMENT STORAGE SCHEME. RANDOM ACCESS AREAS ON THE DRUM AND DISC ARE NOT SHOWN AS THEY WILL REMAIN UNCHANGED.
DR MEANS DOCUMENT REGION.

Apr 1972

- Unified Document System in production
- Further developments/refinements

May 1973

DOCUMENT RETENTION AND ARCHIVAL

It is now possible to retain documents in the document region for a period of up to 99 days. This is done using the SV declaration, either in an EQUIP statement or a MABEL request, appended with nn, where nn is a two-digit number 00-99 indicating the desired retention. For example:

May 1973

A charge of 0.225 system seconds per sector per day of specified retention is made when the SVnn declaration is executed. (Note this differs from the way daily charges are made for documents on the disc document store.)

Brilliant control: charge up-front by size and retention period!
(Terrible for users? But they could specify a short retention period, and extend when needed.)

May 1973

- Automatic Archive
  - done weekly, with two copies if the retention period was long enough
  - Maximum of 99 days retention
- Automatic Retrieval
  - Pre-scan of job, and job-start deferred while retrieving specified documents
  - Retrieval requests “batched at hourly intervals so as to improve the efficiency of retrievals from a number of requests to the same archive tape. The user’s job is run after his document has been retrieved.”
  - Not extended to interactive access
- User Retrieval
  - Access to the archive directory containing all attributes of the document and the tape name and serial number.
  - RETRIEVE statement to restore document to the document region
- Operator-mounted tapes – HSM without tape automation

Jul 1973

- July 1973
  - Cyber 76 installed, taking over main compute role from 3600
- Oct 1973
  - Command to suppress automatic retrieval
- June 1974
  - High priority execution queue for retrievals
- Nov 1975
  - Phase out of 3600 and DR starts
  - Utility to copy files to Cyber 76
- Jul 1976
  - DR charges increase by 80%
- Dec 1976
  - 3600 phase-out begins
- May 1977
  - 3600 service ceased, and document region/HSM died

Dec 1977

- Dec 1977
  - “Perhaps the only answer to the problem of providing a file service lies in an automated archive system. The Division is currently investigating the utility of such devices.”
- Aug 1979
  - First mention of Calcomp Automated Tape library installation
- Aug 1980
  - Article on Braegen ATL (first in Australia), and pictures
  - Facom M190 host
  - XYTEX software to intercept operator mount requests
  - STK round tape drives
  - Robotic tape library
  - NSC Hyperchannel network

ATL

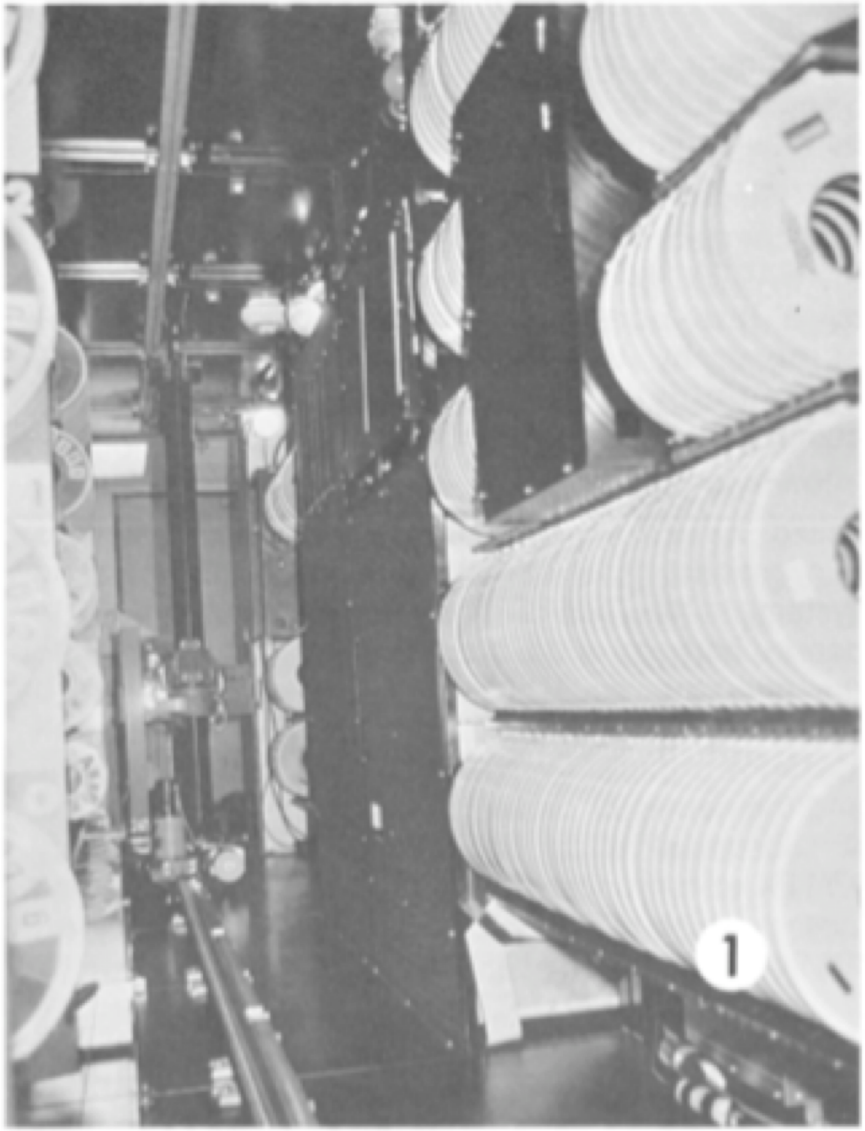

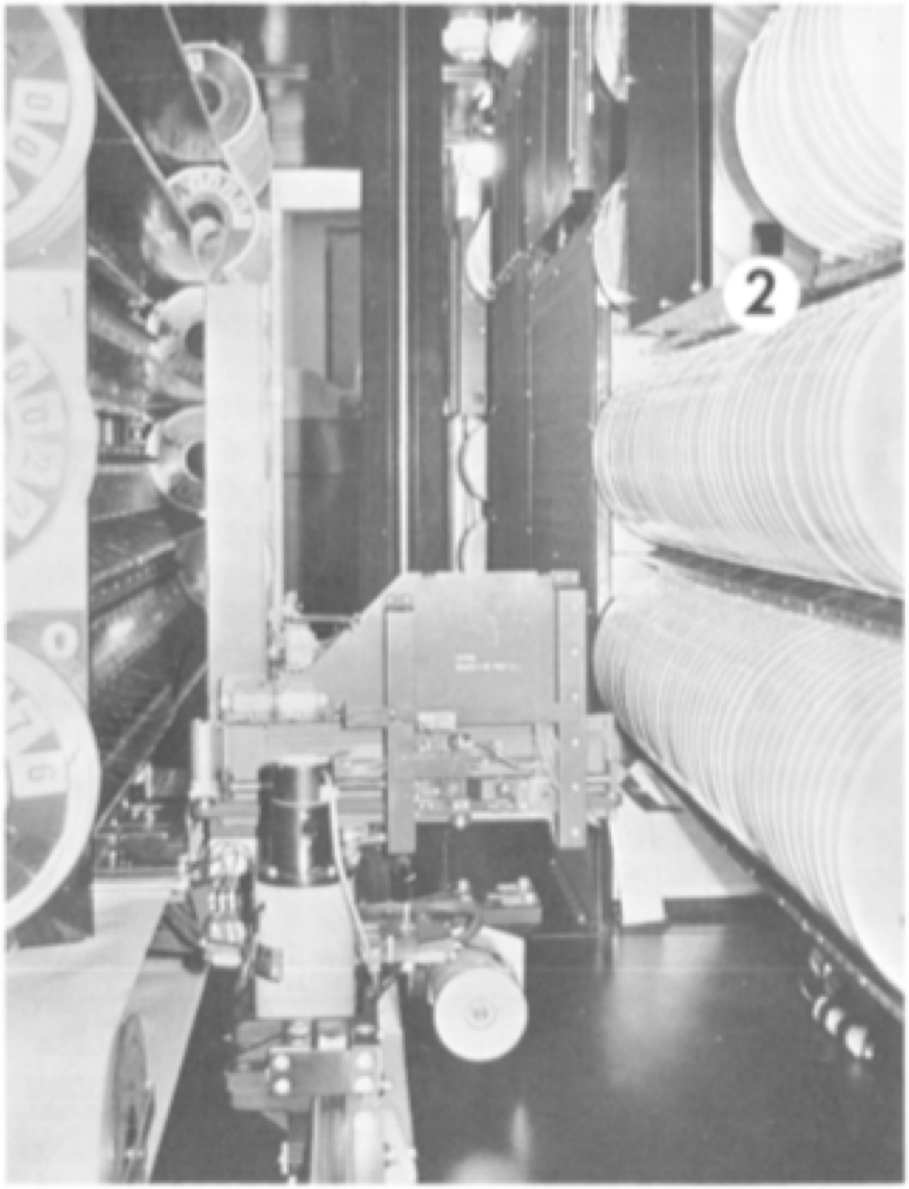

ATL

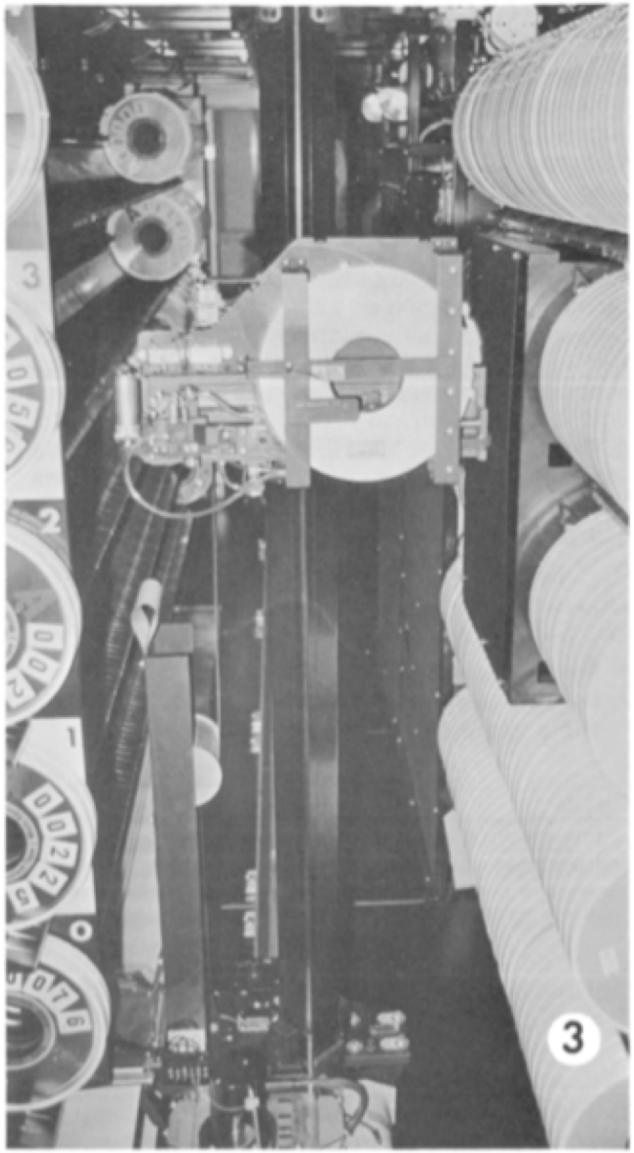

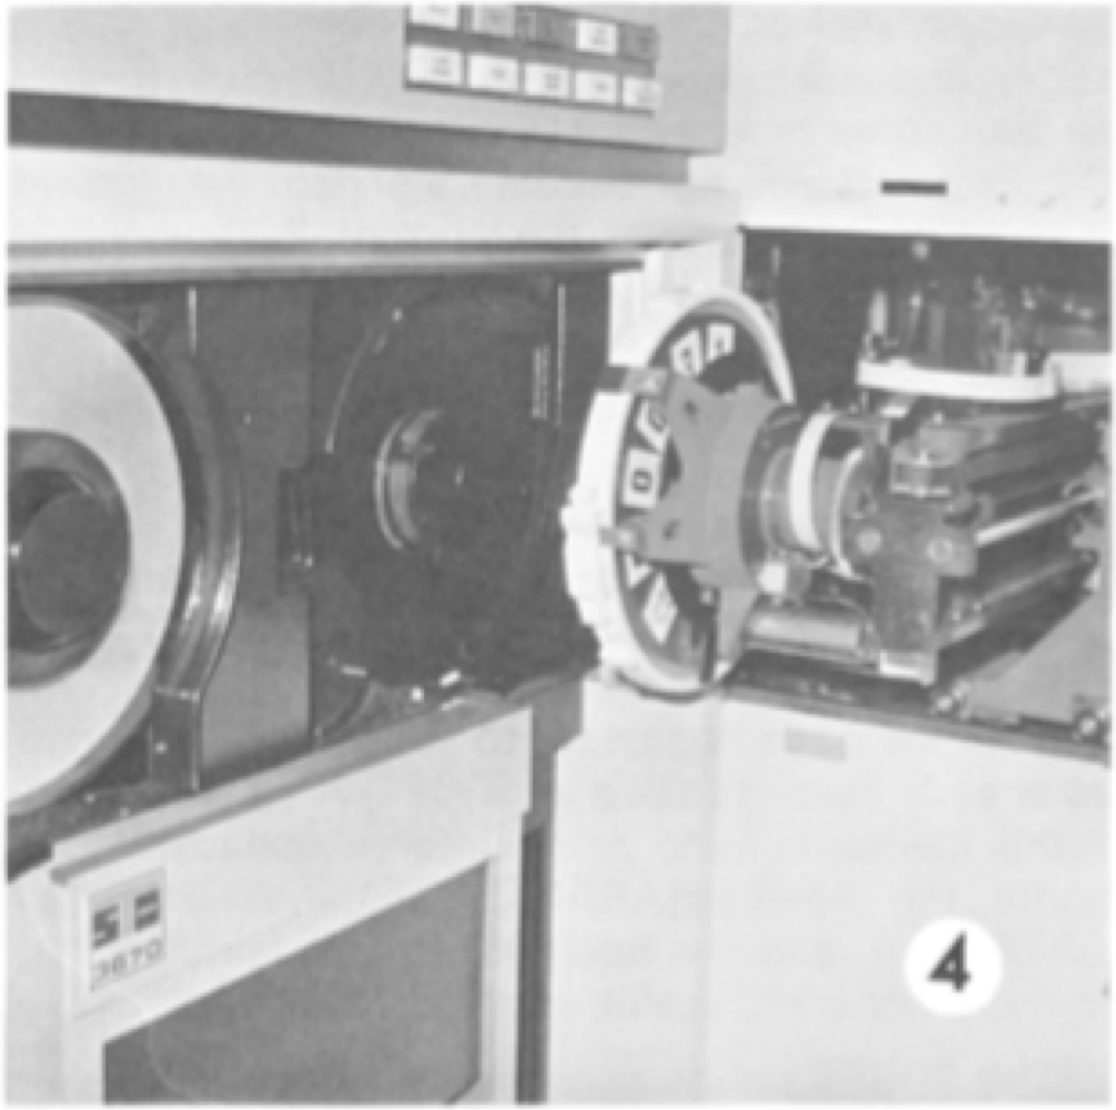

ATL/Terabit File Store

- Feb 1984
  - CSIRO/FACOM developed TFS in production after long development period
  - Archive system
  - Commands to send and retrieve files, audit, change metadata
  - Hierarchical file structure
- Tape automation without HSM

  - 1987-1988 – Csironet sold
  - 1990 – CSIRO ceased using Csironet
  - 1990 – CSIRO HPC restart – Joint Supercomputer Facility – Cray Y-MP

JSF

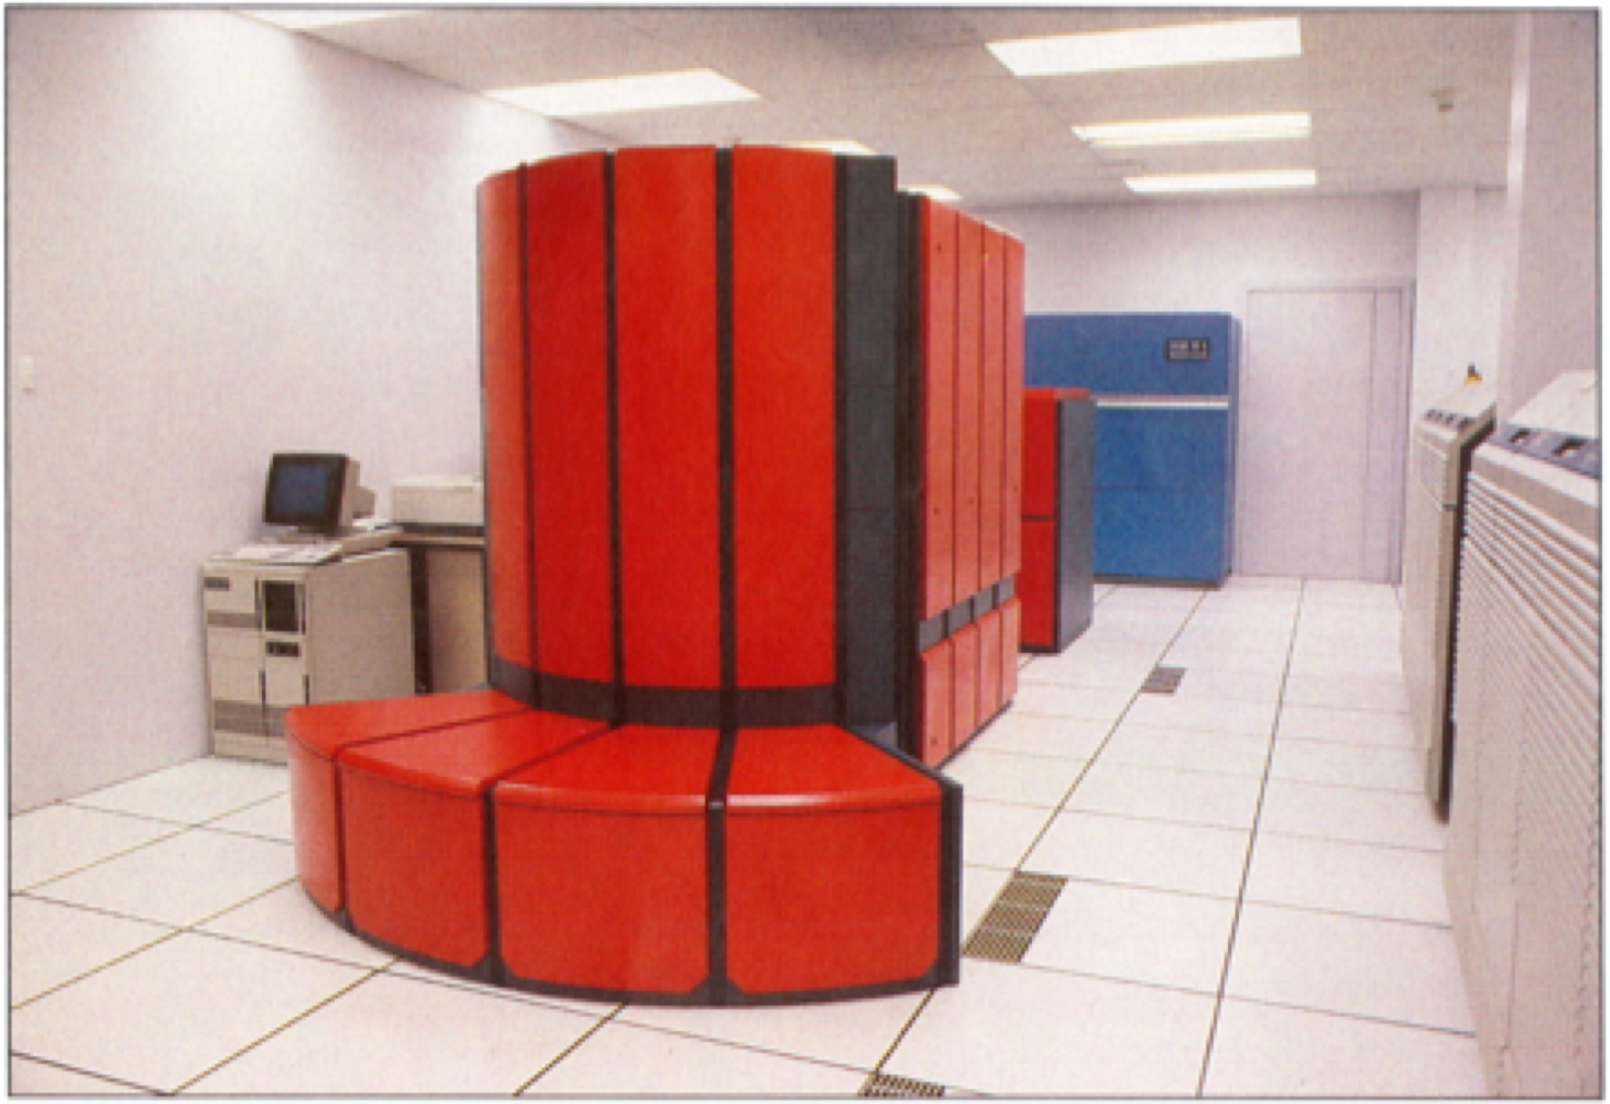

CSIRO DMF history

- 1990 JSF
  - Cray Y-MP turned a compute-bound problem into an i/o and storage problem
- 1991 JSF
  - Searching for solutions for users
    - Interim – users were given access to tapes, ran jobs to tar up files.
  - Information on DMF at CUG
  - Acquisition of STK Cartridge tape drives
  - 14 November – Data Migration Facility started on CSIRO /home area
- HSM without tape automation
- 1992 August: CSF
  - Replacement Cray Y-MP at Uni of Melbourne
- 1993 June: CSF
  - Automation – STK 4400 tape library
- HSM with tape automation for the first time in CSIRO's history

CSIRO DMF history

- Data Store:
  - 5 sites
  - 7 hosts
  - 11 tape technologies
  - 4 principal systems administrators
  - about 7 disc generations
- Other DMF instances in CSIRO since about 2012

Conclusion

- Issues arose from the start of on-line storage, and continue
  - Sharing, quotas, flushing, file numbers, expiry, charging, archiving, backup, HSM
- CSIRO built its own HSM from the mid-1960s to mid-1970s
  - virtual storage system
- CSIRO started DMF in 1991, with automation in 1993
- History – issues for today
  - much the same!

Thank you

CSIRO IMT Scientific Computing
Robert C. Bell
CSIRO HPC National Partnerships
t +61 428 108 333
e Robert.Bell@csiro.au
w www.csiro.au