DTS SOFTWARE, LLC.

Understanding z/OS Disk Metadata VTOCs, VVDSs, Indexes, and More

Steve Pryor
steve@dtssoftware.com
1.919.833.8426 x162

DASD Volume Structures

- Catalog
- VTOC – Volume Table of Contents
  - Locations of datasets and unused space on DASD volumes
- VTOC Index
  - Indexed search for entries in the VTOC, and space maps
- VSAM Volume Data Set (VVDS)
  - Information for VSAM *and non-VSAM (SMS)* datasets
- Other structures
  - FDRABR model DSCB – information for FDRABR

Dataset Basics – Names, Formats, Records

- Data Set = file
  - 44 character DSN=QUAL1.QUAL2.QUAL3…

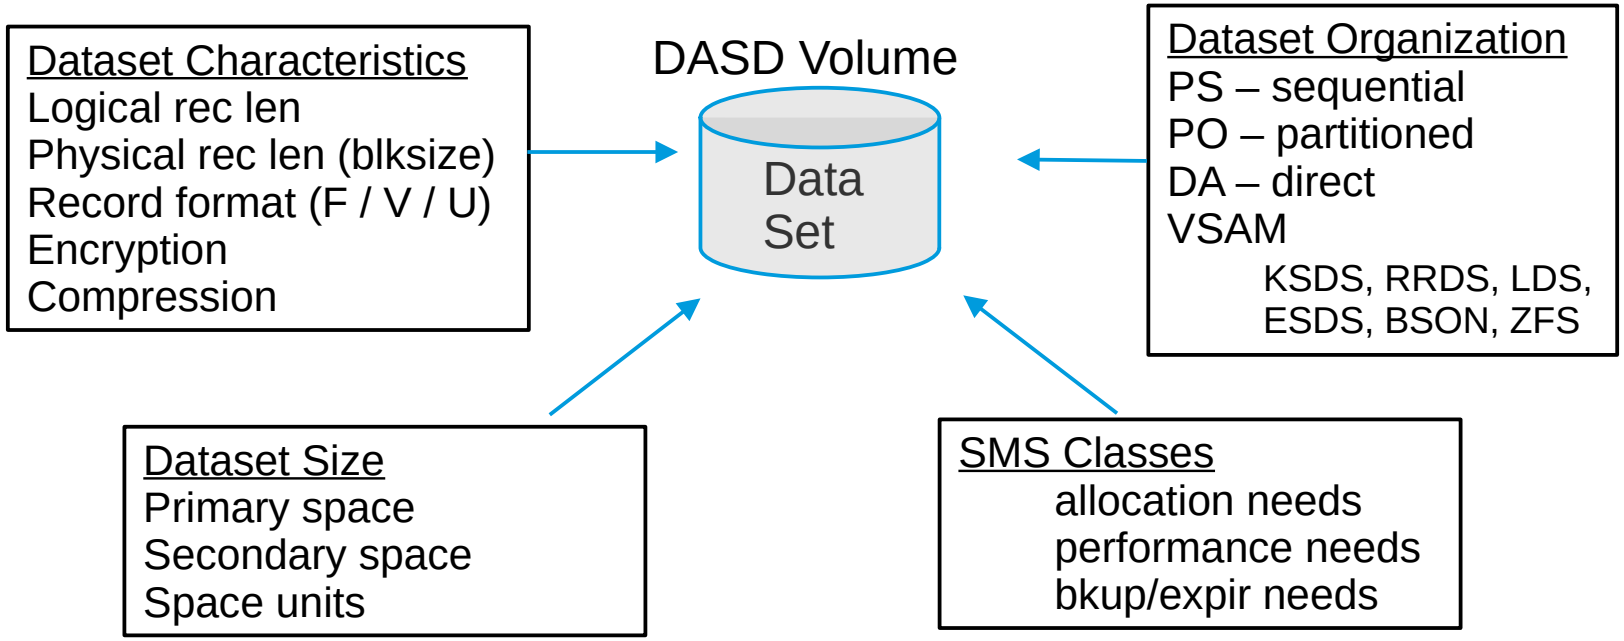

In the Beginning...

1956
IBM 305
Capacity: 5MB
Cost: $3200

1970s - 3330

1970s - 3350

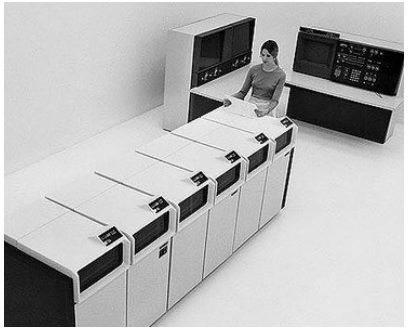

1980s - 3380

1990s - 3390

Platters, Heads, Tracks, and Cylinders

- Traditional Architecture
  - 3390
    - 56664 bytes/track
    - 15 tracks/cylinder
    - Fixed device sizes

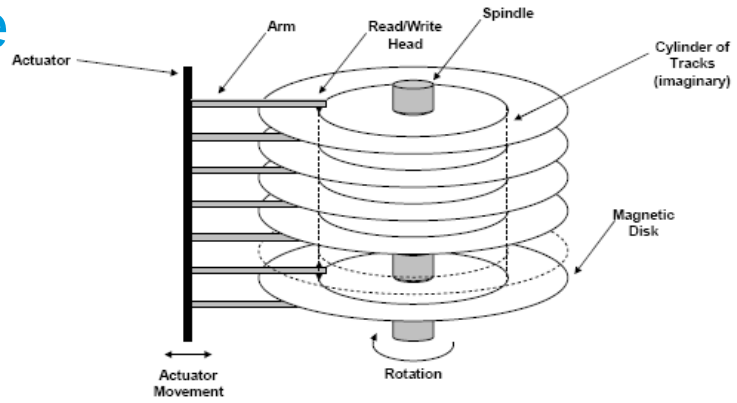

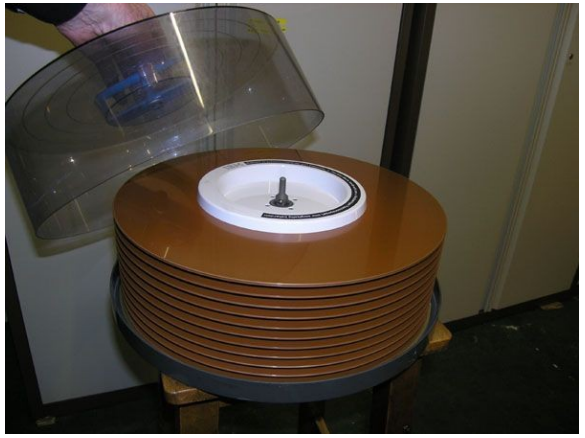

2314 removable pack

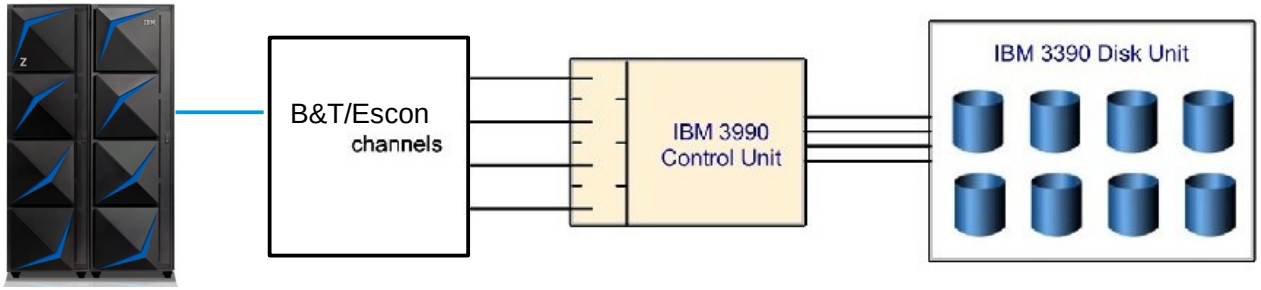

Modern DASD Architecture - RAID

June 1988
Gibson/Katz
UC Berkely
SIGMOD paper

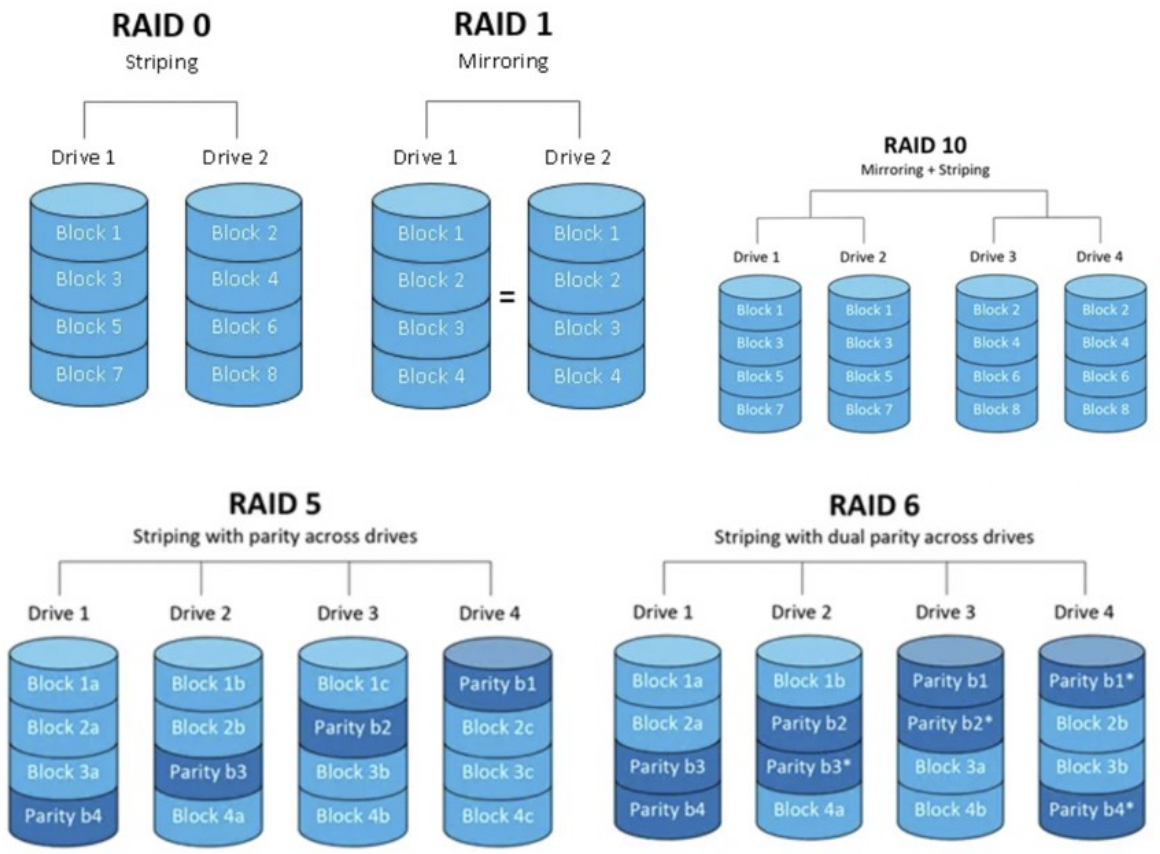

Modern DASD Architecture

- RAID Arrays and Abstraction
- Solid-state memory; storage hierarchy

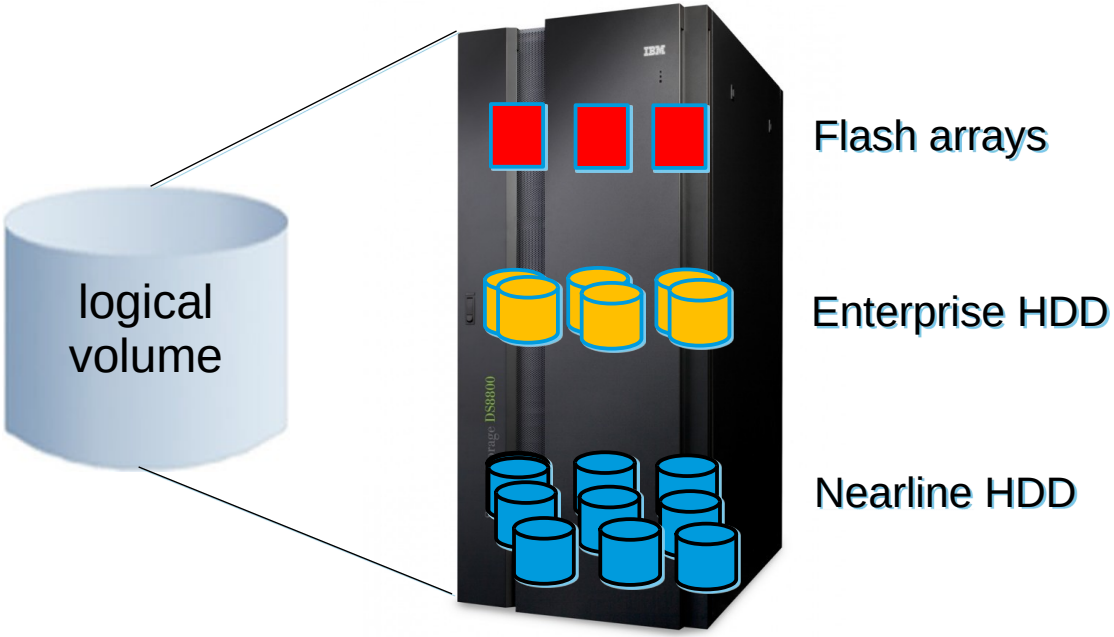

Many complex features

- Adaptive Multistream Prefetch (AMP)
- Easy Tier Integration
- FICON Express
- Flashcopy
- Hi Perf Ficon for z (zHPF)
- HiperPAV
- Intelligent Write Caching (IWC)
- List Prefetch Optimizer
- Multiple Allegiance
- Modified Indirect Address Word (MIDAW)
- Seq Prefetch Adaptive Repl Cache (SARC)
- zHyperWrite

Track Format – Count, Key, Data

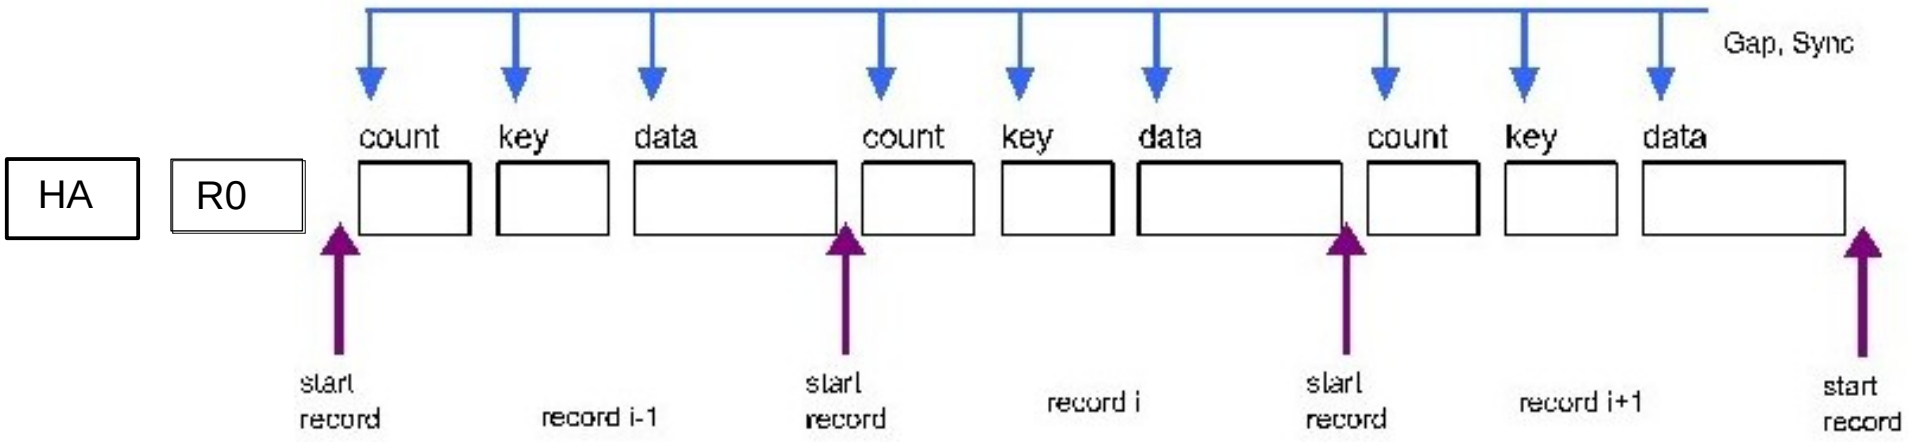

Track address = CCHHR = e.g., 00231E02 = cylinder 35, track 14, record 2

EAV track address = CCCCcccH = e.g., 05DC000E = cylinder 1500, track 14
49F00021 = cylinder 150,000, track 1

Mapping Tracks and Datasets

```
PAGE 0001     5695-DF175  DFSMSDSS V2R04.0 DATA SET SERVICES     2021.297 21:05
 PRINT DATASET(SJP.TEST.DATA) INDYNAM(SMS019)                          00006113
ADR101I (R/I)-RI01 (01), TASKID 001 HAS BEEN ASSIGNED TO COMMAND 'PRINT '
ADR109I (R/I)-RI01 (01), 2021.297 21:05:54 INITIAL SCAN OF USER CONTROL STATEMENTS COMPLETED
ADR016I (001)-PRIME(01), RACF LOGGING OPTION IN EFFECT FOR THIS TASK
ADR006I (001)-STEND(01), 2021.297 21:05:54 EXECUTION BEGINS

*** TRACK(CCHH)  0002000D         R0 DATA  0000000000000000
    COUNT  0002000D010000A0
 0000  E3C8C9E2 40C9E240 D9C5C3D6 D9C440F1  40404040 40404040 40404040 40404040  *THIS.IS.RECORD.1...............*
 0020  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *................................*
 0040  40404040 40404040 40404040 40404040  E3C8C9E2 40C9E240 D9C5C3D6 D9C440F2  *................THIS.IS.RECORD.2*
 0060  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *................................*
 0080  TO  009F   SAME AS ABOVE
    COUNT  0002000D020000A0
 0000  E3C8C9E2 40C9E240 D9C5C3D6 D9C440F3  40404040 40404040 40404040 40404040  *THIS.IS.RECORD.3...............*
 0020  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *................................*
 0040  40404040 40404040 40404040 40404040  E3C8C9E2 40C9E240 D9C5C3D6 D9C440F4  *................THIS.IS.RECORD.4*
 0060  40404040 40404040 40404040 40404040  40404040 40404040 40404040 40404040  *................................*
 0080  TO  009F   SAME AS ABOVE
```

Dataset Extents

- Extent = a group of continguous tracks on a volume
  - May be of varying sizes; need not be in order

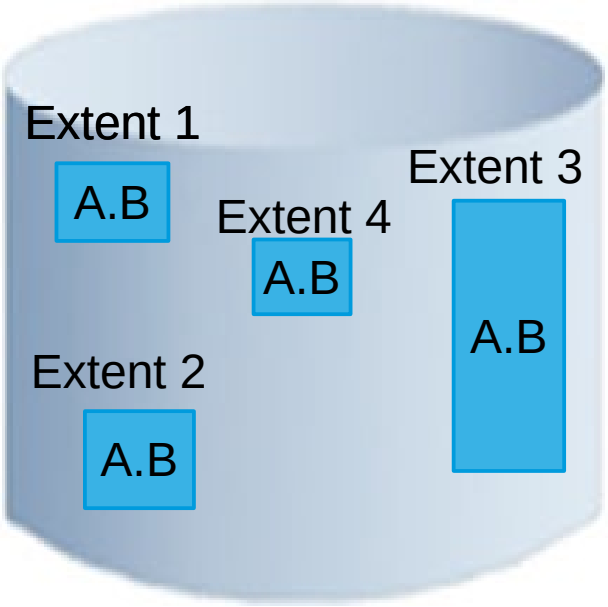

Extent Limits per volume
Sequential = 16
Extended format = 123

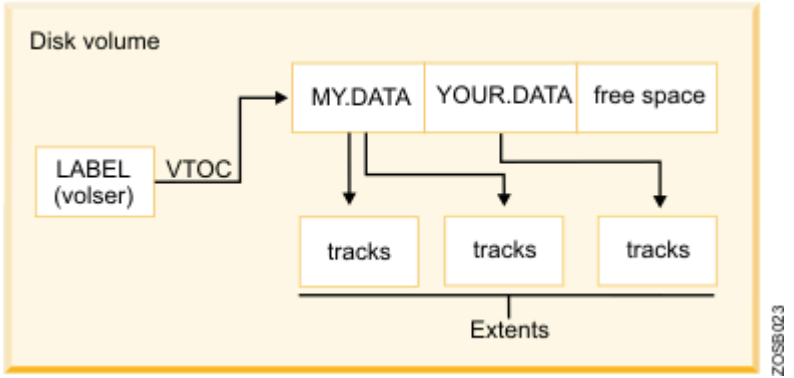

Dataset Extents

```
                CONTENTS OF VTOC ON VOL DTS009  <THIS VOLUME IS NOT SMS MANAGED>
---------------DATA SET NAME----------------  SER NO  SEQNO  DATE.CRE  DATE.EXP  DATE.REF  EXT DSORG RECFM OPTCD BLKSIZE
DTS.TEST71.LOAD                               DTS009      1  2020.013    00.000  2021.297   16  PO    U      00      6144
SMS.IND   LRECL  KEYLEN  INITIAL ALLOC  2ND ALLOC    EXTEND      LAST BLK(T-R-L)  DIR.REM  F2 OR F3(C-H-R)  DSCB(C-H-R)
              0             CYLS              10                      1850   4 46036                0   5   8      0  12   1
EATTR
NS
            EXTENTS  NO  LOW(C-H)   HIGH(C-H)    NO  LOW(C-H)   HIGH(C-H)    NO  LOW(C-H)   HIGH(C-H)
                      0  2571   0     2577 14     1   490   0    499 14      2  1195   0     1204 14
                      3  6698   0     6707 14     4  9825   0    9834 14     5  9953   0     9962 14
                      6  9963   0     9972 14     7  9973   0    9982 14     8 10008   0    10015 14
                      9  9983   0     9984 14    10    82   0      91 14    11   109   0      118 14
                     12   119   0      128 14    13  2628   0    2637 14    14  2638   0     2647 14
                     15  2648   0     2657 14
                                ----ON THE ABOVE DATA SET,THERE ARE    23 EMPTY CYLINDER(S) PLUS  9 EMPTY TRACK(S).
```

Extended Format, Extended Attributes, Extended Addressing

- Extended-format
  - Up to 123 extents / volume
  - SMS-managed
  - Include 32-byte block suffix not visible to applications
    - Used for striping, compression, encryption
- Extended Attributes (EATTR)
  - Dataset allocated in cylinder-managed space
- Extended Addressing
  - VSAM datasets larger than 4G
  - Extent Constraint relief – VSAM datasets more than 255 extents

The ICF Catalog

- Catalog = BCS + VVDS
  - BCS – record types A, B, C,.. X describe datasets, aliases, paths, more
  - VVDS – dataset SYS1.VVDS.Vvolser describes DSNs, components, extents

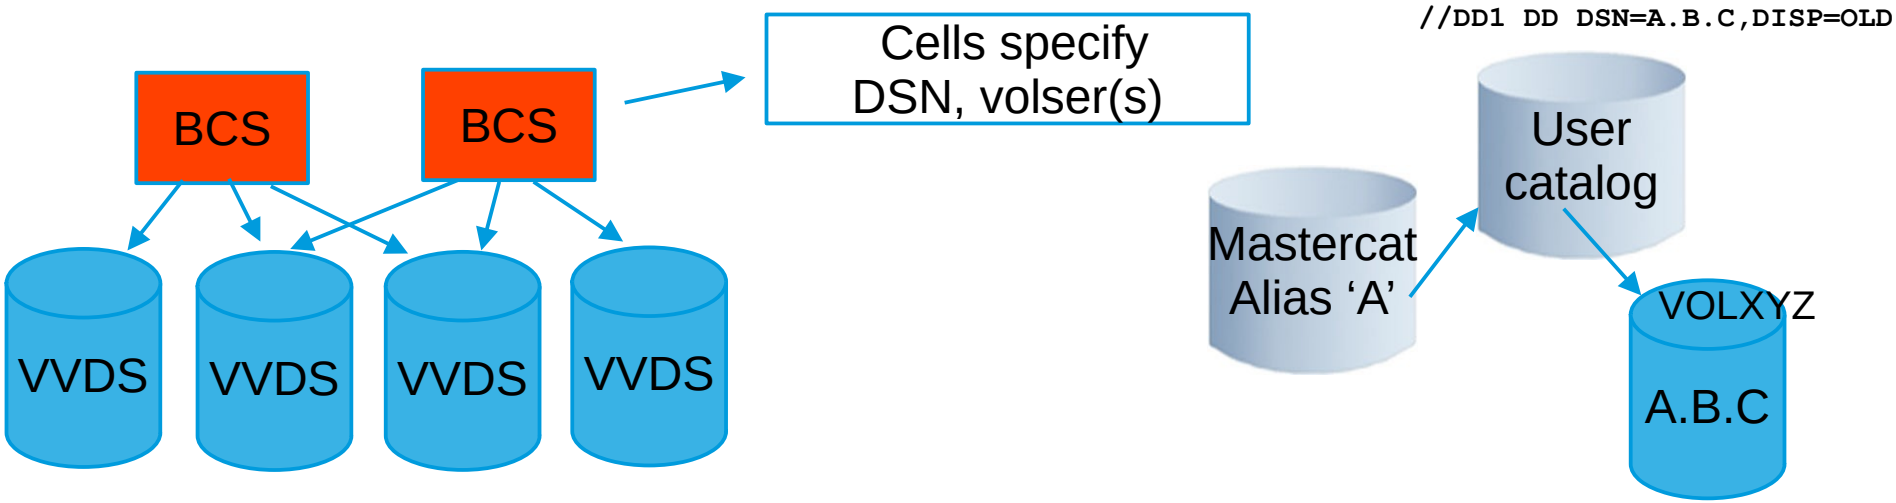

The ICF Catalog

- Locating a data set

```
  LISTC ENT(STEVE.TEST.DS1) ALL
NONVSAM ------- STEVE.TEST.DS1
     IN-CAT --- SJP.USER.CATALOG
     HISTORY
       DATASET-OWNER-----(NULL)      CREATION--------2020.150
       RELEASE----------------2      EXPIRATION------0000.000
       ACCOUNT-INFO-----------------------------------(NULL)
     SMSDATA
       STORAGECLASS -------STC1      MANAGEMENTCLASS-----MAN1
       DATACLASS ---------DATA2      LBACKUP ---0000.000.0000
     ENCRYPTIONDATA
       DATA SET ENCRYPTION-----(NO)
     VOLUMES
       VOLSER------------SMS004      DEVTYPE------X'3010200F'      FSEQN------------------0
       VOLSER------------SMS003      DEVTYPE------X'3010200F'      FSEQN------------------0
     ASSOCIATIONS--------(NULL)
```

Volume Layout

- Track 0
  - Contains the volume label
  - points to the VTOC
- The VTOC
  - Contains DSCBs
  - Points to data sets and free space

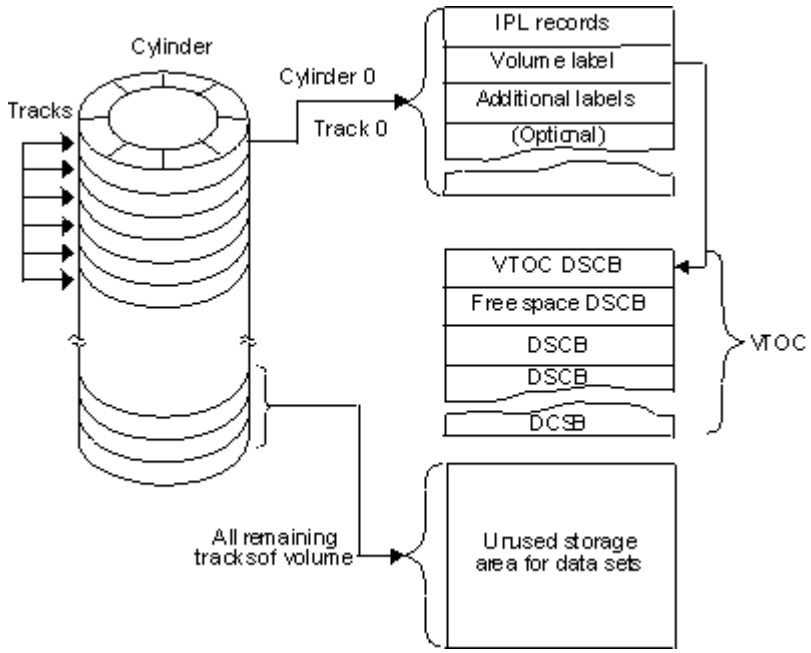

Dataset Information – The VTOC

- 140-char records; 44-character key, 96-byte data
- Dataset Control Blocks (DSCBs)
  - Format 0 – unused DSCB
  - Format 1 – DSN, attributes, points to first 3 extents
  - Format 3 – extents 3+ (F3s are chained together)
  - Format 4 – describes VTOC and volume
  - Formats 5 and 7 – free space
  - Formats 8 and 9 – DSN, attributes, extents for EATTR datasets
  - Formats 2 and 6 – no longer used

Describing the Dataset – F1DSCB / F8DSCB

- Dataset name (key field)
- DSORG, LRECL / BLKSIZE / RECFM
- Primary / secondary space amount
- EXPDT, LRDATE
- Flags
  - SMS, large, encrypted, extended-format, PDSE,
- CCHH for:
  - First 3 extents
  - Chained F3 or F9 DSCB

Finding Datasets - The VTOC Index

- DSN = SYS1.VTOCIX.Vvolser
  - F1DSCB in VTOC points to the VTOC index
- Indexed search improves performance
  - VIER – locations of F1, F4, F8 DSCBs
  - VPSM – free / used space on volume
- Required on SMS-managed volumes
- Created/Formatted/Enabled/Disabled via ICKDSF

More Dataset Information – The VVDS

- VSAM Volume Data Set
  - Required for SMS
    - SMS class names for each dataset
  - Catalog information
    - Related BCS
  - VSAM cluster information
    - Relate cluster name to components
    - RECORG, AVGREC, SPACE, RBAs, buffering, et. al.
    - Component extent descriptors

Mapping the Volume

START TRACK	END TRACK	EXTENT LENGTH	D A T A S E T N A M E	EXTENT	D/S ORG	SPACE ALLOC	2ND. ALLOC	LAST BLK TTR(HEX)	TOTAL TRKS.	EMPTY TRKS.	START CC-HH	END CC-HH
00000	00000	1	*** IPL AND LABEL RECORDS ***								0000-00	0000-00
00001	00029	29	*** VTOC ***	01/01							0000-01	0001-14
00030	00032	3	SYS1.VTOCIX.SMS019	01/01	PS	TRK	0	000215	3	0	0002-00	0002-02
00033	00042	10	SYS1.VVDS.VSMS019	01/01	EF	TRK	10	000901	10	0	0002-03	0002-12
00043	00043	1	SJP.TEST.DATA	01/01	PS	TRK	0	000005	1	0	0002-13	0002-13
00044	00045	2	SJP.INC.X	01/16	PS	TRK	2	001F0F	32	0	0002-14	0003-00
00046	00047	2	SJP.INC.X	02/16							0003-01	0003-02
00048	00049	2	SJP.INC.X	03/16							0003-03	0003-04
00050	00051	2	SJP.INC.X	04/16							0003-05	0003-06
00052	00053	2	SJP.INC.X	05/16							0003-07	0003-08
00054	00055	2	SJP.INC.X	06/16							0003-09	0003-10
00056	00057	2	SJP.INC.X	07/16							0003-11	0003-12
00058	00059	2	*** FREE SPACE ***								0003-13	0003-14
00060	00074	15	STEVE.SJP.TEST.B1	01/01	**	CYL	0	000000	15	0	0004-00	0004-14
00075	00089	15	TOMV.X.DATA	01/01	EF	CYL	1	000E0F	15	0	0005-00	0005-14
00090	00102	13	SJP.PERSH.GDG.G0003V00	01/01	PS	TRK	6591	000102	13	11	0006-00	0006-12
00103	00104	2	SJP.INC.X	08/16							0006-13	0006-14

More Volume Information – FDRABR.Vvolser

- ABR Model DSCB
  - 0-track dataset (F1DSCB only)
  - Backup information and volume flags for FDRABR
    - Current backup generation / cycle / expdt
    - Archive / scratch flags / thresholds
  - Created/updated via PGM=FDRARCH

Summary

- DASD Hardware View

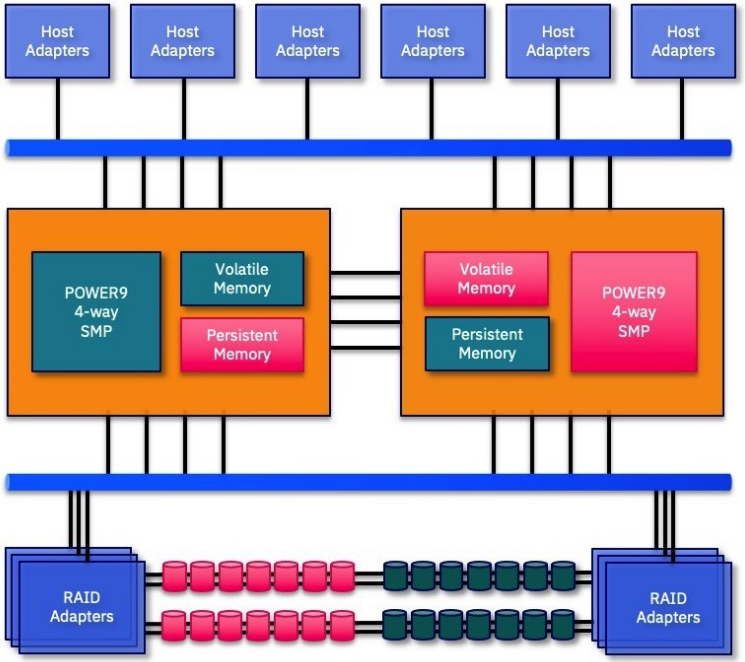

- z/OS View

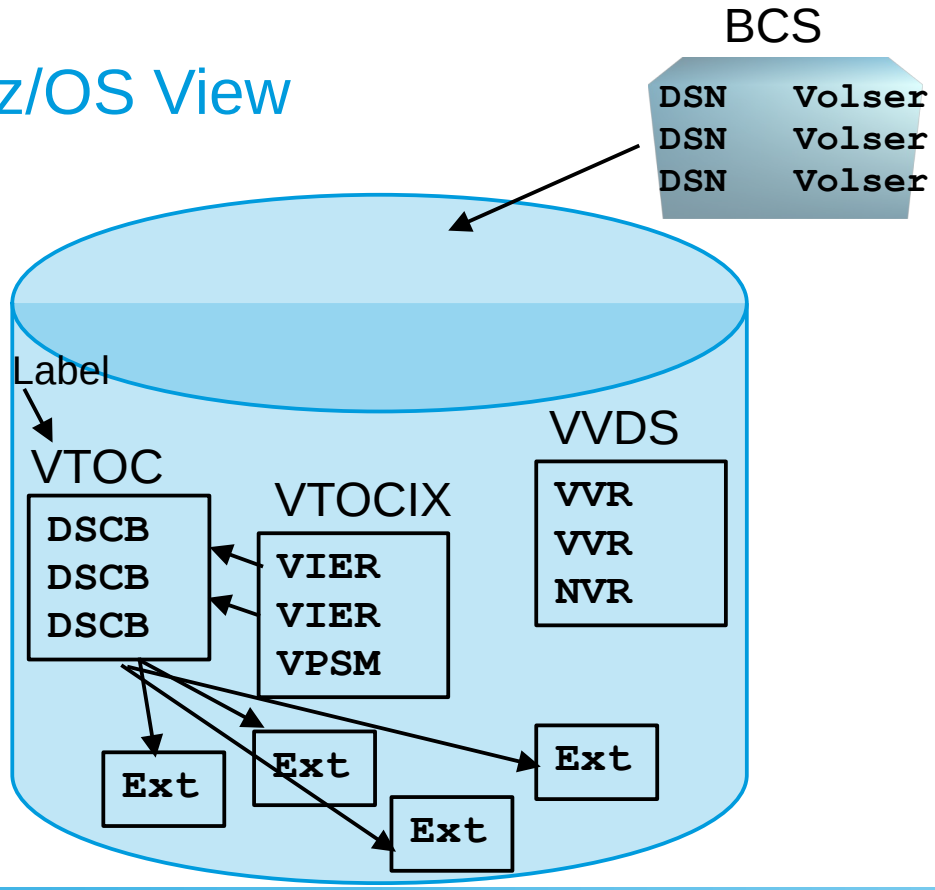

But Wait, There's More...

- Volume Initialization and Replication
  - DVE, FLASHCOPY, PPRC, XRC (Metro / Global Mirror)
- DFSMS
  - Storage Groups, constructs, ACS routines, volume selection...
  - OAM, Tape
- I/O and Control Blocks
  - Access methods
  - UCB, TIOT, DEB, IOB, TIOT, and more
- Backup and Recovery
  - DFSMShsm, FDRABR, CA-DISK

Documentation

- ABCs of z/OS System Programming, Vol 3 SC24-6893
- DFSMS Managing Catalogs SC23-6853
- DFSMSdfp Advanced Services SC23-6861
- DFSMS Using Data Sets SC23-6855

Summary / Q and A

- Next Webinar Nov. 30, 2021
  - *We've Got a Problem! An Introduction to z/OS Dumps and Debugging Tools*

- Reminders
  - DTS Products Release 7.1 available; required for z/OS 2.5
  - New product use available NC for a year – just ask!
  - Send us your ACS routines or DTS product rules for free analysis