

PJ46681 – Support for DASD with more than 65,520 cylinders

2023 TPF Users Group Conference

April 24 – 26 Dallas, TX

Systems Control Program

—

Michael Shershin



Problem Statement

z/TPF online DASD support allows records to be allocated up to 65,520 cylinders. If business growth requires additional records to be allocated and all 65,520 cylinders are used, the only alternative to providing the additional records is to increase the number of DASD modules.

Hardware support to use more than 65,520 cylinders has existed for many years.

Pain Points

Adding DASD modules to increase capacity is painful and a lot of work

- Need a complex outage to increase the number of modules
- Might need to run multiple ODBRs to spread the records across the modules
- Need pool reallocation to add pool records

As-Is User Story

Additional DASD records are needed. Before PJ46681, the following options are available:

- Add records on existing DASD modules as long as the number of cylinders used is less than 65,520
- Add more DASD modules to z/TPF

To-Be User Story

Support up to 1,182,006 cylinders per module in z/TPF

- Online DASD modules only
- Support allocation of all sizes of pool and fixed records (381 bytes, 1055 bytes, and 4 KB)

Space limitation on existing DASD modules is increased eighteenfold.

- No need to add modules to z/TPF due to space limitations on DASD modules.

Technical Details

The following does not support more than 65,520 cylinders

- 3380 and 3390 DASD modules in 3380 emulation mode do not support more than 65,520 cylinders
- General files do not support more than 65,520 cylinders
- General data sets, except for VSAM, do not support more than 65,520 cylinders
- Vertically allocated records

Technical Details

MCHR address is a 14-digit token that is the physical location of a record. Prior to PJ46681 the format is:

- MMMMCCCCHHHRR
 - MMMM: 4-digit hexadecimal number is the symbolic module number
 - CCCC: 4-digit hexadecimal number that is the cylinder number
 - HHHH: 4-digit hexadecimal number that is the head number (or track number)
 - Note: only 1 digit is needed for the head. Therefore, the value is 000H.
 - RR: 2-digit hexadecimal number that is the record number

Technical Details

MCHR address with PJ46681 the format is:

- MMMMCCCCcccHRR
 - MMMM: 4-digit hexadecimal number is the symbolic module number
 - CCCC: 4-digit hexadecimal number that is the low order 16 bits of the cylinder number
 - ccc: 3-digit hexadecimal number that is the high order 12 bits of the cylinder number
 - H: 1-digit hexadecimal number that is the head number (or track number)
 - RR: 2-digit hexadecimal number that is the record number

Technical Details

Example of MCHR address with a cylinder above 65,520 cylinders

```
==> WP/ZMCHR 00001F000040A4BC
CSMP0097I 11.58.27 CPU-B SS-WP    SSU-WP1  IS-01
MCHR0002I 11.58.27 CURRENT      ALTERNATE  RECORD    RECORD
                MMCCHHR          MMCCHHR      SIZE      DUPLICATION
                -----          -----      -----  -----
                001E11D7001701          4K          DUPE
END OF DISPLAY+
```

```
==> WP/ZDFIL 001E11D7001701
CSMP0097I 11.58.45 CPU-B SS-WP    SSU-WP1  IS-01
DFIL0011I 11.58.45 BEGIN DISPLAY OF FILE ADDRESS 001E11D7001701
 00000000- 00000000 E3C5E2E3 00000000 00000000 .....TEST .....
 00000010- 00000000 00000000 00000000 00000000 .....
END OF DISPLAY - ZEROED LINES NOT DISPLAYED+
```

Technical Details

Multi-path lock facility (MPLF) in DASD control units is used by loosely coupled z/TPF complexes to hold record locks. z/TPF uses an 8-byte lock format. Prior to PJ46681, the format is:

- CCHRF00V
 - CC: 2-byte hexadecimal number is the cylinder number
 - H: 1-byte hexadecimal number is the head number (or track number)
 - R: 1-byte hexadecimal number that is the record number
 - F: 1-byte flag byte that contains flags and the low order 6 bits of the SDA
 - 00 – Reserved and the value is set to zero
 - V – VFA indicators and a user indicator

Technical Details

The MPLF lock format that z/TPF uses with PJ46681 is:

- CCHRFccV
 - CC: 2-byte hexadecimal number is [the low order 16 bits of the cylinder number](#)
 - H: 1-byte hexadecimal number is the head number (or track number)
 - R: 1-byte hexadecimal number that is the record number
 - F: 1-byte flag byte that contains flags and the low order 6-bits of the SDA
 - [cc – 2-byte hexadecimal number that is the high order 12 bits of the cylinder number](#)
 - V – VFA indicators and a user indicator

Technical Details

How do you tell z/TPF that more than 65,520 cylinders are being used?

- In SIP the ONLFIL macro defines the number of cylinders that a device type can use.
- ONLFIL DEVICE_x=(3390, number_of_cylinders)
- For 3390, the maximum number of cylinders allowed is increased to 1,182,006

Technical Details

Example of SIP ONLFIL macro with more than 65,520 cylinders

```
*****
*                ONLFIL - ONLINE DEVICE CHARACTERISTICS                *
*****
ONLFIL      DUPMP=Y,      DUP MODULE PAIRING STATUS (Y=prime dupe)      X
            LIRECOVERY=Y,      LOST INTRPT RECOVERY                    X
            LOSTINT=10,      LOST INTRPT TIMER                        X
            DEVICEA=(3390,70119), TYPE OF DEVICE  NUM OF CYCLS        X
            DUPTYPA=S,      DUPLICATION STATUS OF DEVICE            X
            PERMA=12,      NBR OF PERMANENTLY MOUNTED MODULES        X
            IPLABLE=6,      NBR OF IPLABLE MODULES                  X
            VSNCHAR=WX,      ALPHA PORTION OF VSN                    X
            NAMDEVA=WPXX,      X                                      X
            VOLNOA=1,      STARTING NUMERIC                          X
            DEVICEB=(3390,10017), TYPE OF DEVICE                    X
            NAMDEVB=DEVB,      X                                      X
            DUPTYPB=S,      DUPLICATION STATUS OF DEVICE            X
            PERMB=20,      NBR OF PERMANENTLY MOUNTED MODULES        X
            VOLNOB=13,      STARTING VOLUME SERIAL NUMBER            X
            CTK6EXP=Y
```

Technical Details

Module configuration table (CTSD)

- Includes all cylinders
- Online formatter
- Module copy uses the CTSD
 - ZMCPY ALL
 - ZMCPY UP
- Capture and Restore uses the CTSD

Technical Details

Display of the module configuration table (command ZDMCT)

==> WP/ZDMCT

CSMP0097I 19.18.45 CPU-B SS-WP SSU-WP1 IS-01

DMCT0003I 17.18.45 DISPLAY OF ACTIVE MODULE CONFIGURATION TABLE

TRACK SIZE RANGES FOR DEVA

INDEX	FROM		TO		TOTAL	SIZE	DUPE	
	CYLINDER	TRK	CYLINDER	TRK	TRACKS			
1	0	02	1	06	20	LARGE	YES	—
2	1	07	1	07	1	LARGE	NO	
...								
386	62286	03	63286	02	15000	4K	YES	
387	63286	03	64286	02	15000	4K	YES	
388	64286	03	65286	02	15000	4K	YES	
389	65286	03	66286	02	15000	4K	YES	
390	66286	03	67286	02	15000	4K	YES	—
391	67286	03	67385	09	1492	4K	YES	
	67385	10	70103	06	40767	UNDEFINED		
392	70103	07	70118	14	233	4K	YES	

END OF DISPLAY+

Technical Details

Display of allocated pool segments (command ZPOOL DISPLAY)

==> WP/ZPOOL DISPLAY OPM-AAA SECTION-4D6A

CSMP0097I 19.20.12 CPU-B SS-WP SSU-WP1 IS-01

POOL0313I 17.20.12 POOL SECTION DETAIL - OPM-AAA PST-CUR

PST CREATED - 02/15/2023 17:06:28 PST RECORD TYPE - #PSTV1

CYL	STARTING			FIRST MOD	NBR MOD	POOL		BASE ORDINAL	RECORDS ALLOCATED
	H	R	MOD			SECT	RCC		
592	13	1	0000	0000	6	4D6A	94	000000000000000000	1800
594	8	1	0000	0000	6	4D6A	94	000000000000000708	1800
596	3	1	0000	0000	6	4D6A	94	000000000000000E10	3600
599	8	1	0000	0000	6	4D6A	94	00000000000001C20	3600
...									
655	13	1	0000	0000	6	4D6A	94	00000000000109C8	360
656	3	1	0000	0000	6	4D6A	94	0000000000010B30	360
656	8	1	0000	0000	6	4D6A	94	0000000000010C98	360
65504	7	1	0000	0000	6	4D6A	94	0000000002BA48A	2031666
70103	7	1	0000	0000	6	4D6A	94	0000000004AA4BC	16776

END OF DISPLAY+

Value Statement

When additional DASD records are needed, support for up to 1,182,006 cylinders provides the option to allocate new records on existing DASD modules

- No z/TPF outage is needed
- Faster to implement than adding more modules
 - Increasing DASD space can be done in weeks rather than months
- Less risk because updates are only needed to the FCTB and a pool reallocation
- Assumes that access rate to DASD modules after expansion is not an issue
- Archive more data on z/TPF

Conclusion

APAR PJ46681 provides this support. It is available. It was shipped in December 2022.

Thank you

© Copyright IBM Corporation 2022. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available at [Copyright and trademark information](#).

