What's New in z/VM I/O?

Eric Farman (farman@us.ibm.com)

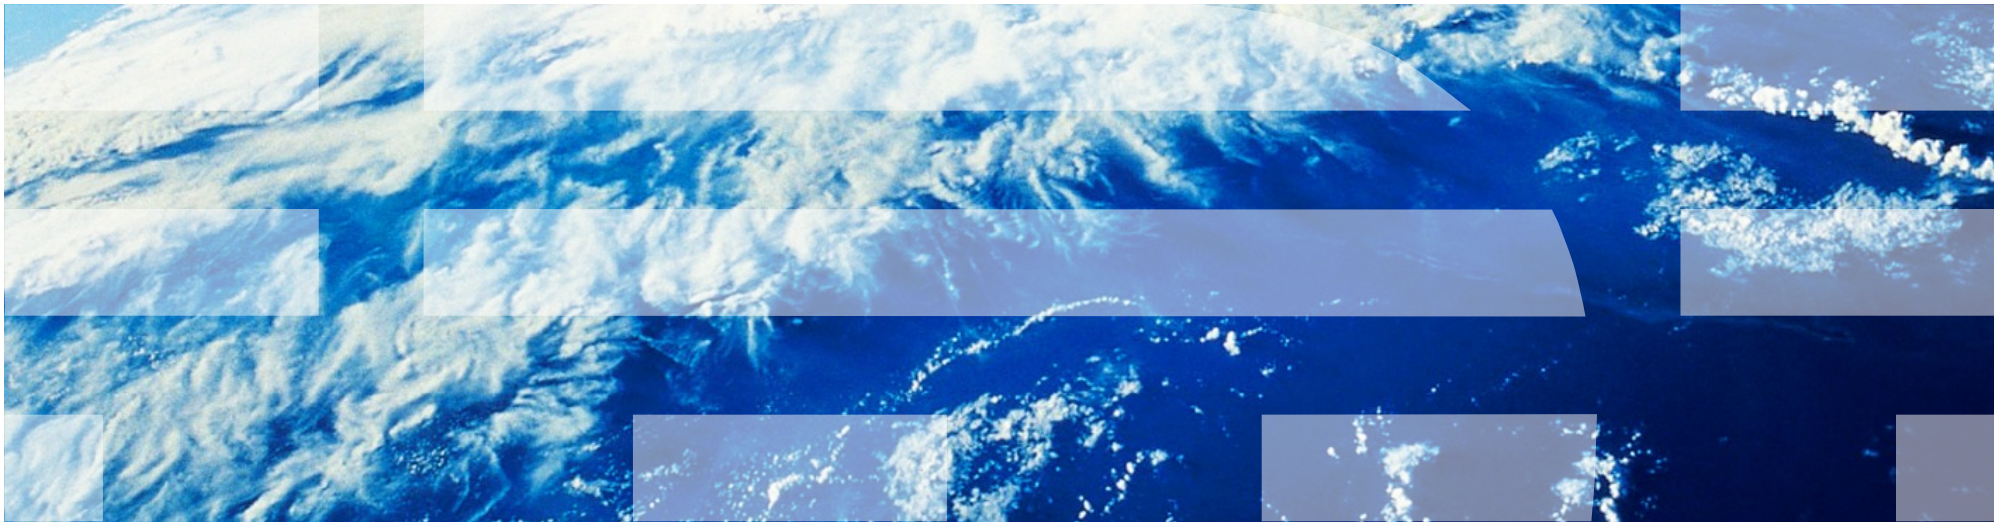

What's New in z/VM I/O?

- Extended Address Volumes (EAV)
- FlashCopy / Space-Efficient
- Query DASD DETAILS

Extended Address Volumes (EAV)

- Two APARs for z/VM 5.4.0 and 6.1.0
  - VM64709 (CP)
  - VM64711 (CMS)
- PTFs available December, 2009

Extended Address Volumes (EAV)

- Mechanism for creating and using volumes larger than 65,520 cylinders (the so-called "mod-54" variant of a 3390-9)
  - Identified as 3390-A devices
- Cylinder addressing is done by "stealing" 12 bits from head value of CCHH (ccccchhhh)
  - Often documented as ccccCCCh, where the complete cylinder is rearranged as CCCccccc

Extended Address Volumes (EAV)

- EAV volumes can be attached to the SYSTEM, generally for the purposes of minidisks
- Non-PERM extents on CPOWNED EAV volumes are restricted to the first 65,520 cylinders
  - Extents that cross this line will be truncated, with message HCP138E
  - Extents that exist entirely above this line will be ignored, with message HCP139E
- CPFMTXA will enforce this boundary requirement

Extended Address Volumes (EAV)

- z/VM support for dedicated devices, and fullpack minidisks
  - "Fullpack" is defined as 0-END, or DEVNO
  - The ending cylinder must be "END", even if the number would equal the size of the volume
- Partial pack minidisks on 3390-A volumes are supported, provided they exist completely below cylinder 65,520

Extended Address Volumes (EAV)

- Diagnose xA8 will operate on any area of disk, provided the application is aware of EAV addressing constructs
  - New SGIOPTS byte, immediately after SGILPM, is added, with flag bit SGIEAV (x80) defined to signal this awareness
  - Any application attempting to reference above cylinder 65,520 without enabling this bit will fail with CC=1 R15=2 (unsupported device)

Extended Address Volumes (EAV)

```
*** SGIOP - Synchronous General I/O Parameters
*
*     +-------------+------+------+--------------------------+
*   0 |  SGIDEVNO   |SGIKEY|SGIFLG|         SGIRESV1         |
*     +-------------+------+------+--------------------------+
*   8 |          SGICPA           |         SGIRESV2         |
*     +---------------------------+------+------+------------+
*  10 |         SGICCWA           |:DEVST|:SCHST|  SGIRESCT  |
*     +------+------+-------------+------+------+------------+
*  18 |SGILPM|:OPTS |  SGIRESV3   |  SGIRESV4   |  SGISNSCT  |
*     +------+------+-------------+-------------+------------+
*  20 |         SGIRESV5          |         SGIRESV6         |
*     +---------------------------+--------------------------+
*  28 |         SGIRESV7          |         SGIRESV8         |
*     +---------------------------+--------------------------+
*  30 |         SGIRESV9          |         SGIRESVA         |
*     +---------------------------+--------------------------+
*  38 |                                                      |
*     =                      SGISDATA                        =
*     |                                                      |
*     +------------------------------------------------------+
*  58
*
*** SGIOP - Synchronous General I/O Parameters
```

Extended Address Volumes (EAV)

- Diagnose x210 returns the size of a volume in field VRDCPRIM, but this field cannot hold anything larger than 16-bits (65,535)
  - Field practical maximum is 65,520 (xFFF0)
- If VRDCPRIM contains xFFFE, indicates the field has overflowed (large, EAV volume)
- New 32-bit field VRDCCYLS at offset x4C
  - Always contains cylinder size

Extended Address Volumes (EAV)

```
*** VRDCBLOK - VIRTUAL/REAL DEVICE CHARACTERISTICS BLOCK
*
*     +-------------+-------------+------+------+------+------+
*   0 |  VRDCDVNO   |  VRDCLEN    |:CVCLA|:CVTYP|:CVSTA|:CVFLA|
*     +------+------+------+------+------+------+------+------+
*   8 |:CRCCL|:CCRTY|:CCRMD|:CCRFT|:CUNDV|:CRDAF|  VRDCRSVD   |
*     +------+------+------+------+------+------+-------------+
*  10 |  VRDCCUTY   |:CCUMD|  VRDCDVTY   |:CDVMD| VRDCDVFE-   |
*     +------+------+------+------+------+------+-------------+
*  18 |-(016)|:CSDFE|:CDVCL|:CDVCO|  VRDCPRIM   |  VRDCTRKC   |
*     +------+------+------+------+-------------+------+------+
*  20 |:CSECT|     VRDCTOTR       |   VRDCHA    |:CMODE|:CMDFR|
*     +------+------+-------------+-------------+------+------+
*  28 |  VRDCNKOV   |  VRDCKOVH   |  VRDCALTC   |  VRDCALTR   |
*     +-------------+-------------+-------------+-------------+
*  30 |  VRDCDIG    |  VRDCDIGN   |  VRDCDVCY   |  VRDCDVTR   |
*     +------+------+------+------+-------------+-------------+
*  38 |:CMDR |:COBR |:CCUID|////////////////////////////////////|
*     +------+------+------+////////////////////////////////////|
*     |////////////////////////////////////////////////////////|
*     +------+--------------------+--------------------------+
*  48 |:RCUC |////////////////////|         VRDCCYLS         |
*     +------+--------------------+--------------------------+
*  50 |                       VRDCPGID                        |
*     |                    +-----------------------------------+
*  58 |                    |///////////////////////////////////|
*     +--------------------+-----------------------------------+
*  60
*
*** VRDCBLOK - VIRTUAL/REAL DEVICE CHARACTERISTICS BLOCK
```

Extended Address Volumes (EAV)

- Other Diagnoses (x18, x20, xA4, x250, and *BLOCKIO) cannot operate on devices that exist wholly or partially above cylinder 65,520
  - Attempts to do so will be rejected, with the proper "unsupported device" return/condition code combination for the Diagnose

Extended Address Volumes (EAV)

- New DDR fully supports EAV volumes
- Updates to FlashCopy ensure it fully supports EAV volumes
- New fields in MRSEKSEK (Domain 7, Record 1) monitor record less likely to wrap for larger volumes
  - CALCURCY, CALSKCYL, IORPOSSM, CALECYL

Extended Address Volumes (EAV)

- CMS currently supports volumes up to 32,767 cylinders in size
- With aforementioned APAR, CMS is updated to support, via FORMAT and ACCESS, volumes and minidisks up to 65,520 cylinders in size
  - Since this does not provide support for EAV-sized volumes, can be applied separately from CP APAR

FlashCopy / Space-Efficient

- Space-Efficient volumes are those created solely for the purpose of being targets persistent FlashCopy relationships
  - A special LIC feature of DS8000
- Tracks for the volume are only allocated as needed to maintain the point-in-time image as source/volume change

FlashCopy / Space-Efficient

- Persistent FlashCopy relationships are used to maintain pointers between source and target volumes, and changes made to either one
- The no-background-copy option (typically NOCOPY) is used to prevent tracks being written to target device until they are actually needed by storage subsystem.

FlashCopy / Space-Efficient

- z/VM APAR VM64449 (December, 2008) was released to add persistent FlashCopy, and by extension, FlashCopy/SE, support
  - Several APARs released later; check for related!
- Several new commands are added to create and manage these relationships

FlashCopy / Space-Efficient

- FLASHCOPY ESTABLISH
  - Create Persistent Relationships
  - Needed for SE volumes, but can be used with traditional volumes, if desired
  - REVERSIBLE option is combination of CHGRECORD and NOTGTWRITE, and ensures targets is not SE
  - LABEL and SAVELABEL options also exist on traditional (non-persistent) FLASHCOPY command

FlashCopy / Space-Efficient

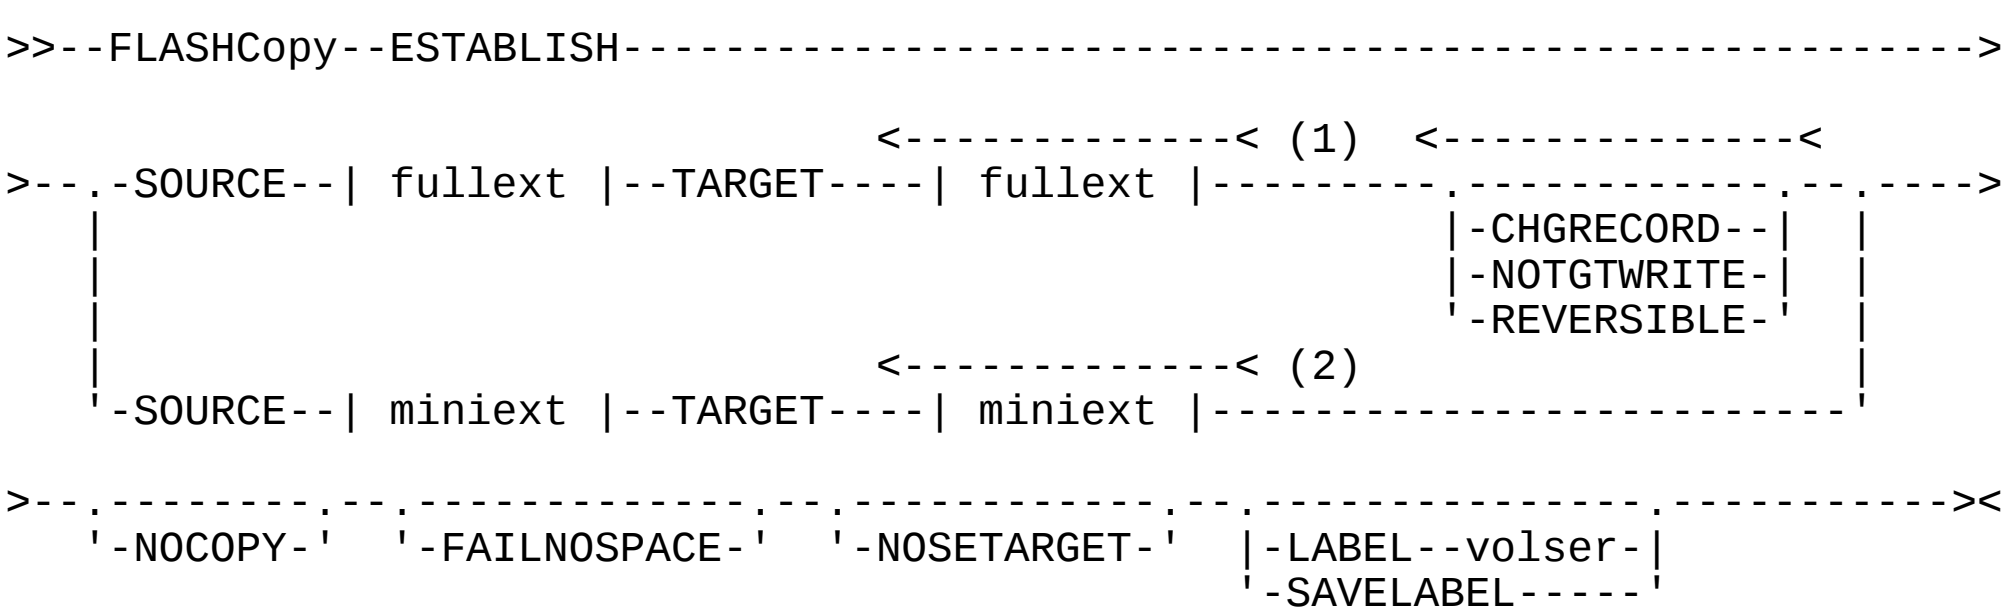

FlashCopy / Space-Efficient

- FLASHCOPY WITHDRAW
  - Removes Persistent FlashCopy relationship between paired volumes
  - FORCE option required when background copy has not completed, as a result of NOCOPY option on Establish
    - FLASHCOPY BACKGNDCOPY can be used to initiate background copy without withdrawing relationship, but cannot be used on space-efficient targets

FlashCopy / Space-Efficient

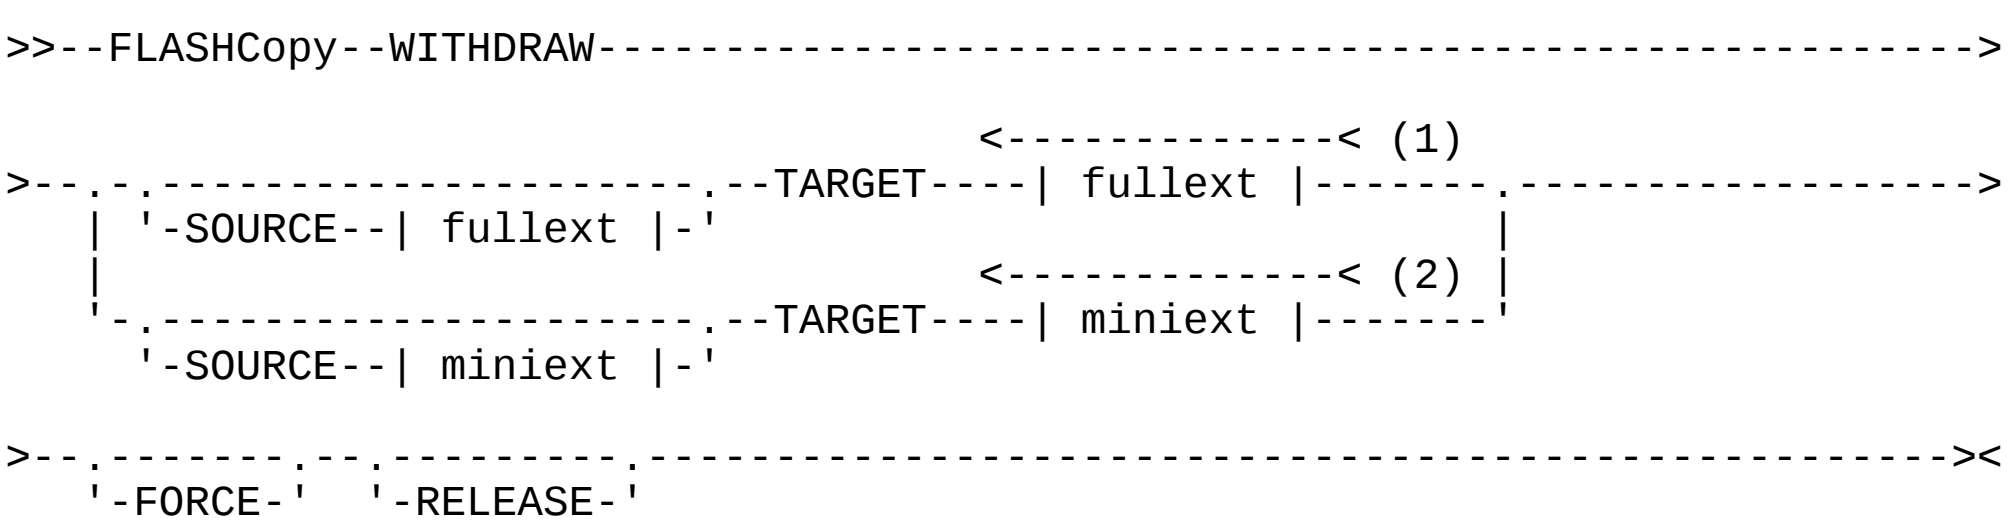

FlashCopy / Space-Efficient

- FLASHCOPY RESYNC
  - For persistent relationships, will establish a new point-in-time copy between source and target volume(s)
- FLASHCOPY TGTWRITE
  - Disables the write-inhibit option on target volume, if specified on original Establish

FlashCopy / Space-Efficient

- QUERY FLASHCOPY
  - Privileged command to interrogate status of FlashCopy relationship(s) on a real device
  - State on the hardware, and options for retrieving information about a relationship created within a z/VM session during current IPL

FlashCopy / Space-Efficient

```
                                <---------------<
>>--Query--FLASHCopy--.-HARDWARE----.-----------.---.-----------------------><
                      |              |-rdev------|   |
                      |              '-rdev-rdev-'   |
                      |-TABLEd-----------------------|
                      |-CREATed--userid|*------------|
                      |-OWNer--userid|*--------------|
                      |-SEQUENCE--hhhhhhhh-----------|
                      |-VOLume--volser---------------|
                      '-DEVice--rdev-----------------'
```

FlashCopy / Space-Efficient

```
q flashcopy hardware
               --------SOURCE-------- --------TARGET--------
SEQUENCE FLGS RDEV VOLSER CC...CC/HH RDEV VOLSER CC...CC/HH REMAINING/TOTAL
4B145C34 0800 5100 PACK01     100/00 5101 PACK02     100/00 0/1500
4B14700A 0800 5100 PACK01     200/00 5101 PACK02     200/00 0/1500
4B145C34 8800 5100 PACK01     100/00 5101 PACK02     100/00 0/1500
4B14700A 8800 5100 PACK01     200/00 5101 PACK02     200/00 0/1500

q flashcopy tabled
SEQUENCE ---DATE--- --TIME--  RDEV  VOLSER CREATOR  OWNER    VDEV
4B145C34 2008-04-30 15:19:55  5100> PACK01 RWS      RWS      0200
4B145C34 2008-04-30 15:19:55 >5101  PACK02 RWS      RWS      0300
4B14700A 2008-04-30 15:20:00  5100> PACK01 RWS      RWS      0210
4B14700A 2008-04-30 15:20:00 >5101  PACK02 RWS      RWS      0310
     Ready; T=0.01/0.01 15:20:40
```

Query DASD DETAILS

- New microcode on DS8000 boxes allow for encrypted DASD and/or solid-state drives
- z/VM APAR VM64650 (May, 2009) provides updates to Query DASD DETAILS output to indicate the presence of these facilities
  - No new output is displayed if neither exists

Query DASD DETAILS

```
q dasd details 521d
521D  CUTYPE = 2107-E8, DEVTYPE = 3390-0A, VOLSER = ERF001, CYLS = 3339
      CACHE DETAILS:  CACHE NVS CFW DFW PINNED CONCOPY
           -SUBSYSTEM   Y    Y   Y   -    N       N
              -DEVICE   Y    -   -   Y    N       N
      DEVICE DETAILS: CCA = 1D, DDC = --, DED = YES, SSD = NO
      DUPLEX DETAILS: --
      CU DETAILS: SSID = 0102, CUNUM = 5200

q dasd details dead
DEAD  CUTYPE = 6310-80, DEVTYPE = 9336-10, VOLSER = ERF105, CYLS = 9446
      BLKS = 7340032
      DEVICE DETAILS: DED = YES, SSD = NO
```

Fin