z9 109 - 2094
MIDA Facility

Kenneth Trowell
zSeries Channel Development
September 9, 2005
kmtrowel@us.ibm.com

IBM System z9 109 - Main storage data movement

- Data can be move from one storage location to another as a result of executing a CP storage movement instruction, or data can be moved to or from storage to an externally connected I/O device as a result of executing an I/O operation (for example to or from a DASD / Disk / device)
- For these data movement operations there is a data address of where to move the data to or from and either a length or count (or other implicit condition) controlling the mount of data to be moved or transferred
- For all data movement cases, explicitly stated or implicitly implied data addressing rules have to be followed, e.g. 2K / 4K addressing boundary rules
- These instruction / I/O operation data addressing rules can be so restrictive to some instructions / I/O that a different set of programming tasks may be used. Use of these different programming tasks my not be efficient in the use of the processing capacity of the element performing the task (i.e. a FICON channel)
- When this occurs either changes or additions are provided to the processing architecture (z/Architecture)
- The introduction of MIDA (change for MIDA) is an example of this

z9 109 - Main storage CP instruction data movement

- zArchitecture CP instructions are used to move data between storage locations or change data in the CPs main storage.
  - Example: Move Character - MVC - instruction

MVC	L	B1	D1	B2	D2

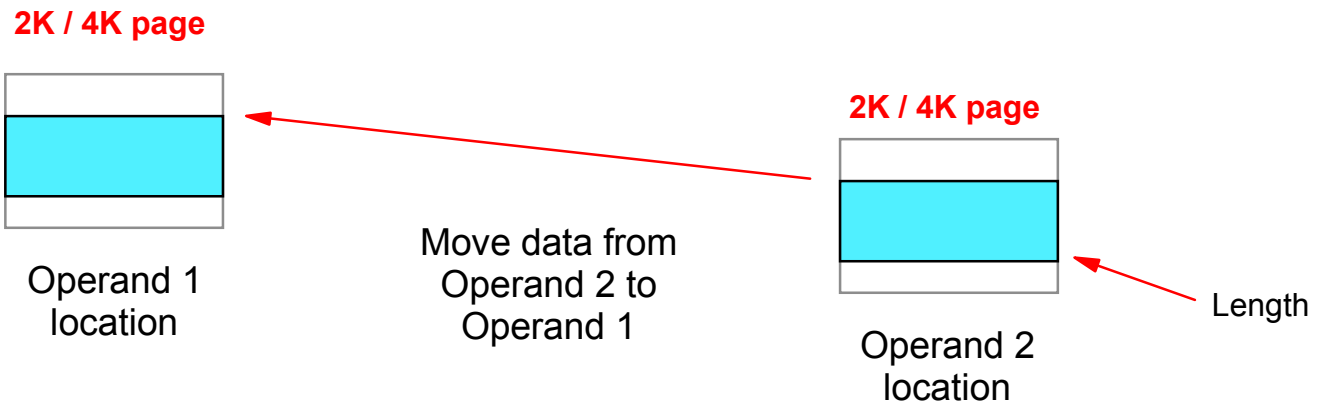

z9 109 - Main storage I/O operation data movement

- To transfer data to or from an I/O device, to or from CP main storage the zArchitecture CP Start Subchannel instruction is used, this provides the channel with the 'channel program' (CCW chain)

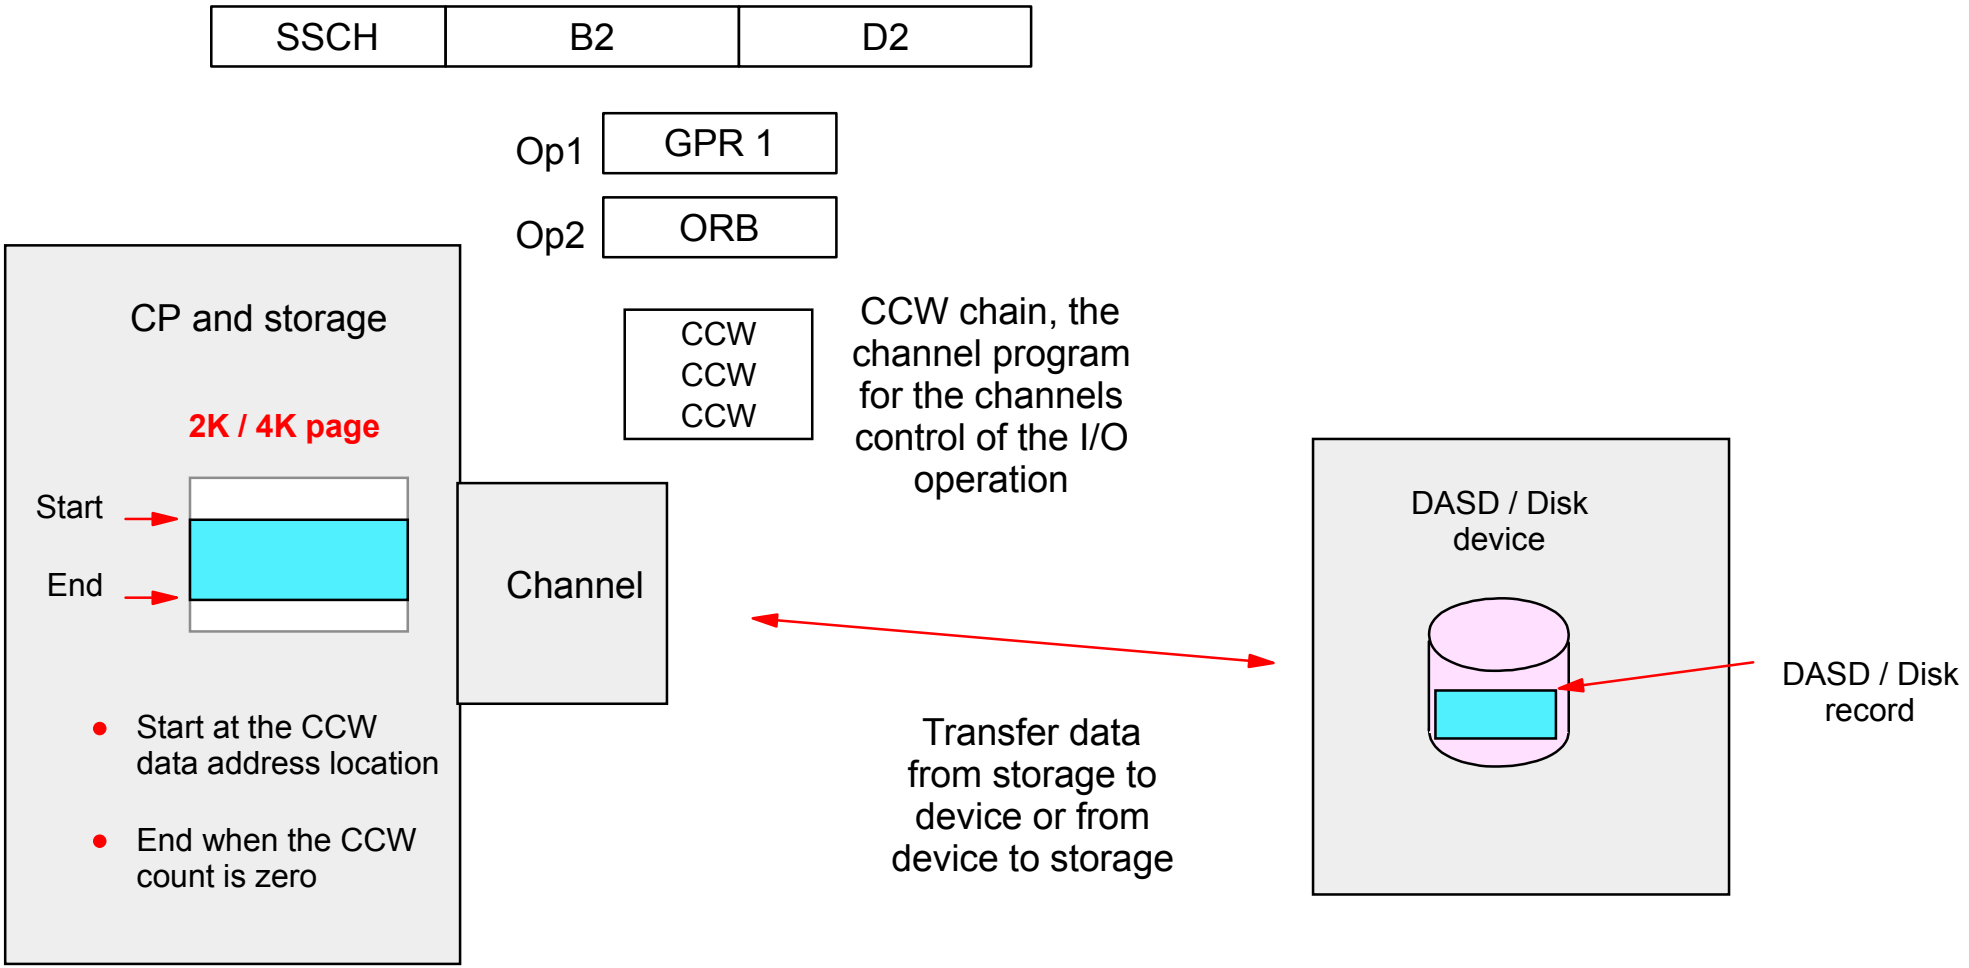

I/O Data Transfer Requirement - 1 I/O operation - Real Addressing)

Generally applications generally do not get contiguous real storage greater than 4k in length (or across a 2K / 2K address boundary). Applications using virtual storage allocations can get greater then 2K / 4K of contiguous virtual storage, but this will be virtual addressing, the allocation will not be in contiguous real. Two things, the channel always uses real storage addressing and for a single I/O operation (CCW data transfer) the data will be sent contiguously as part of the same I/O operation.

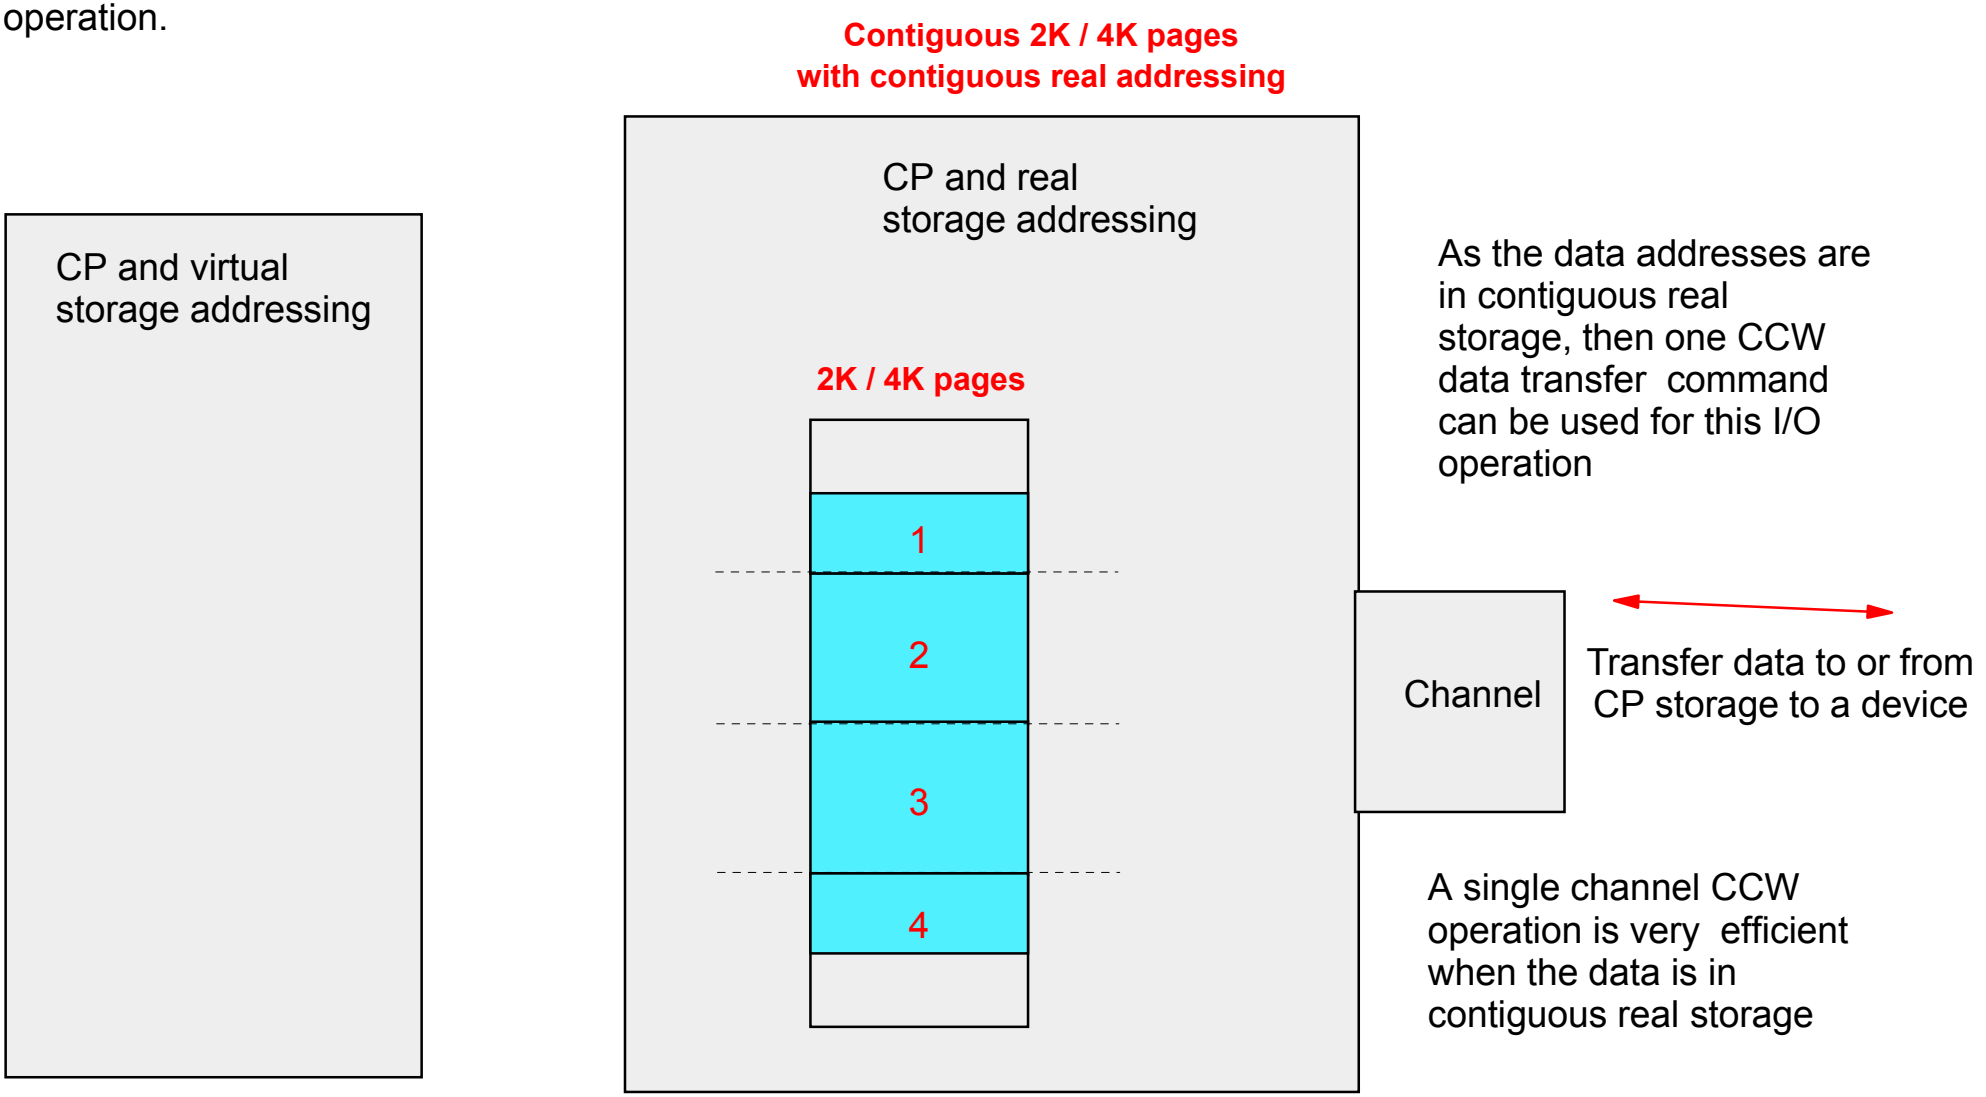

I/O Data Transfer Requirement - 1 I/O operation (IDA)

Generally applications do not get contiguous real storage greater than 4K in length (or across a 2K / 4K address boundary) Applications is using virtual storage allocations can get greater then 2K / 4K of contiguous virtual storage, but this will be virtual addressing, the allocation will not be in contiguous real. Two things, the channel always uses real storage addressing and for a single I/O operation (CCW data transfer) the data will be sent contiguously as part of the same I/O operation.

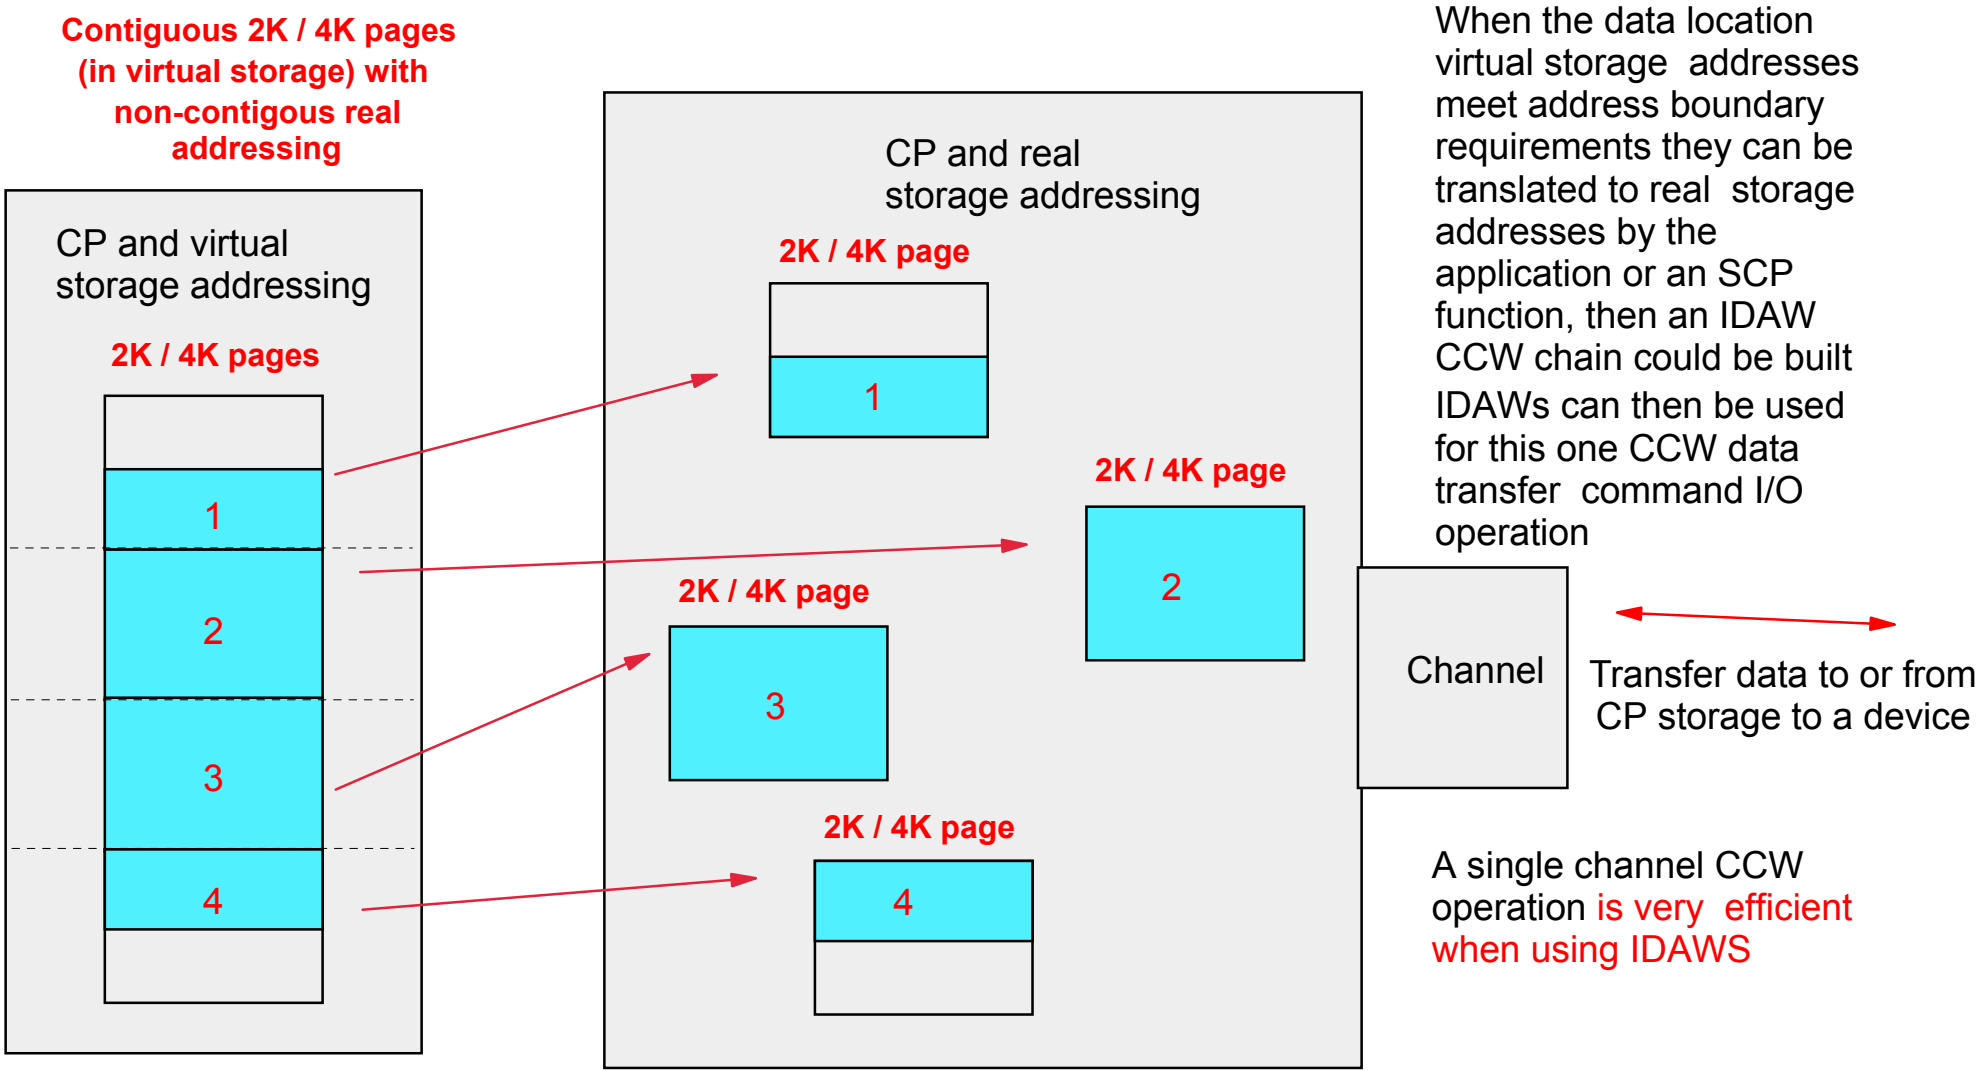

When the data location virtual storage addresses meet address boundary requirements they can be translated to real storage addresses by the application or an SCP function, then an IDAW CCW chain could be built IDAWs can then be used for this one CCW data transfer command I/O operation

A single channel CCW operation is very efficient when using IDAWS

I/O Data Transfer Requirement - 1 SCATTER / GATHER I/O operation (CD)

Some application or subsystem single CCW I/O operations require that data is fetched or stored is in scattered storage locations. These scatter / gather storage address locations and length do not meet the requirements of IDAWS, so IDAWS cannot be used. In this case a Chain Data CCW I/O operation can be used. However the use of Chain data is very inefficient as pseudo CCW channel commands are sent for each chain data CCW.

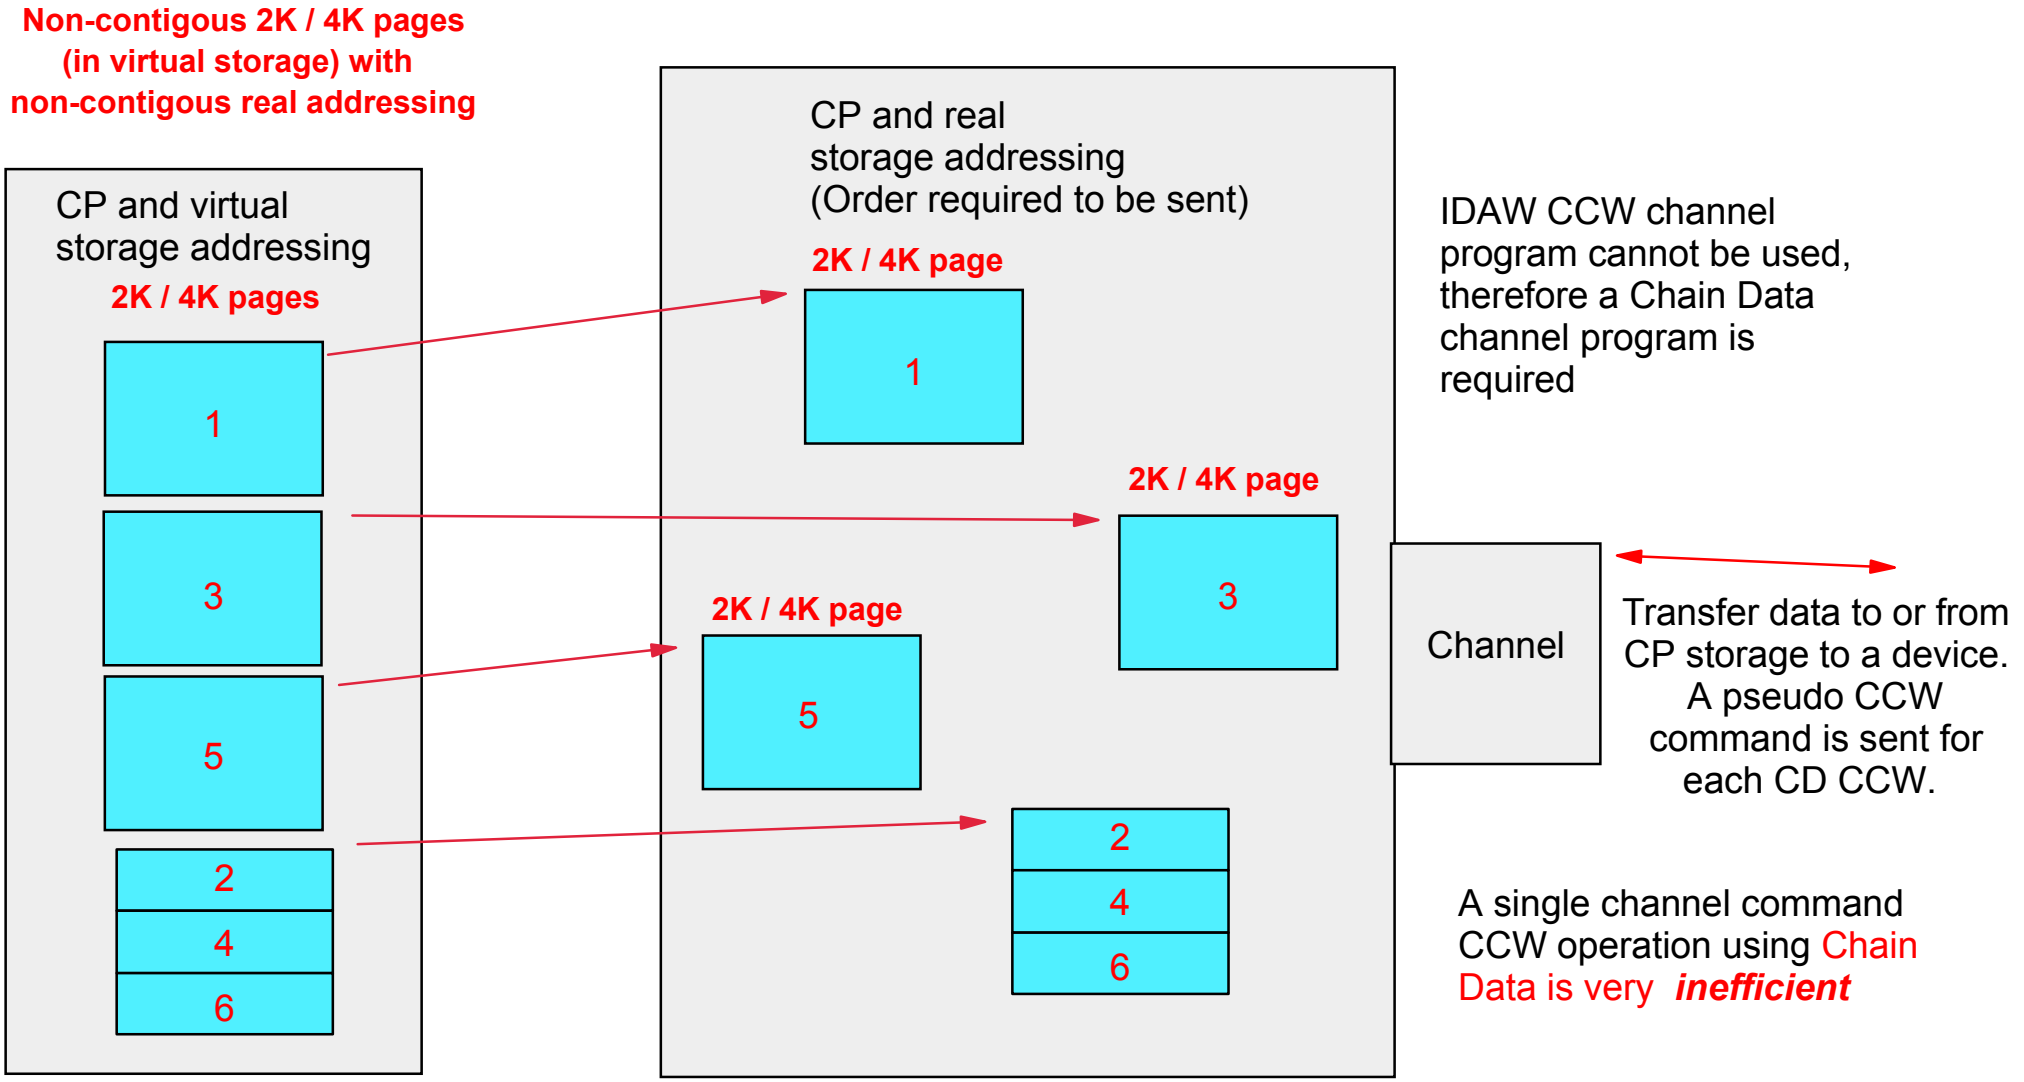

I/O Data Transfer Requirement - 1 SCATTER / GATHER I/O operation (MIDA)

Some application or subsystem single CCW I/O operations require that data is fetched or stored is in scattered storage locations. These scatter / gather storage address locations and length do not meet the requirements of IDAWS, so IDAWS cannot be used. IBM has introduced a new channel operation called **MIDAWS**. Use of MIDAWS has the channel efficiency of an IDAW operation and the flexibility of Chain Data storage location addressing (MIDAWS is supported on selected processors and channels).

Non-contigous 2K / 4K pages (in virtual storage) with non-contigous real addressing

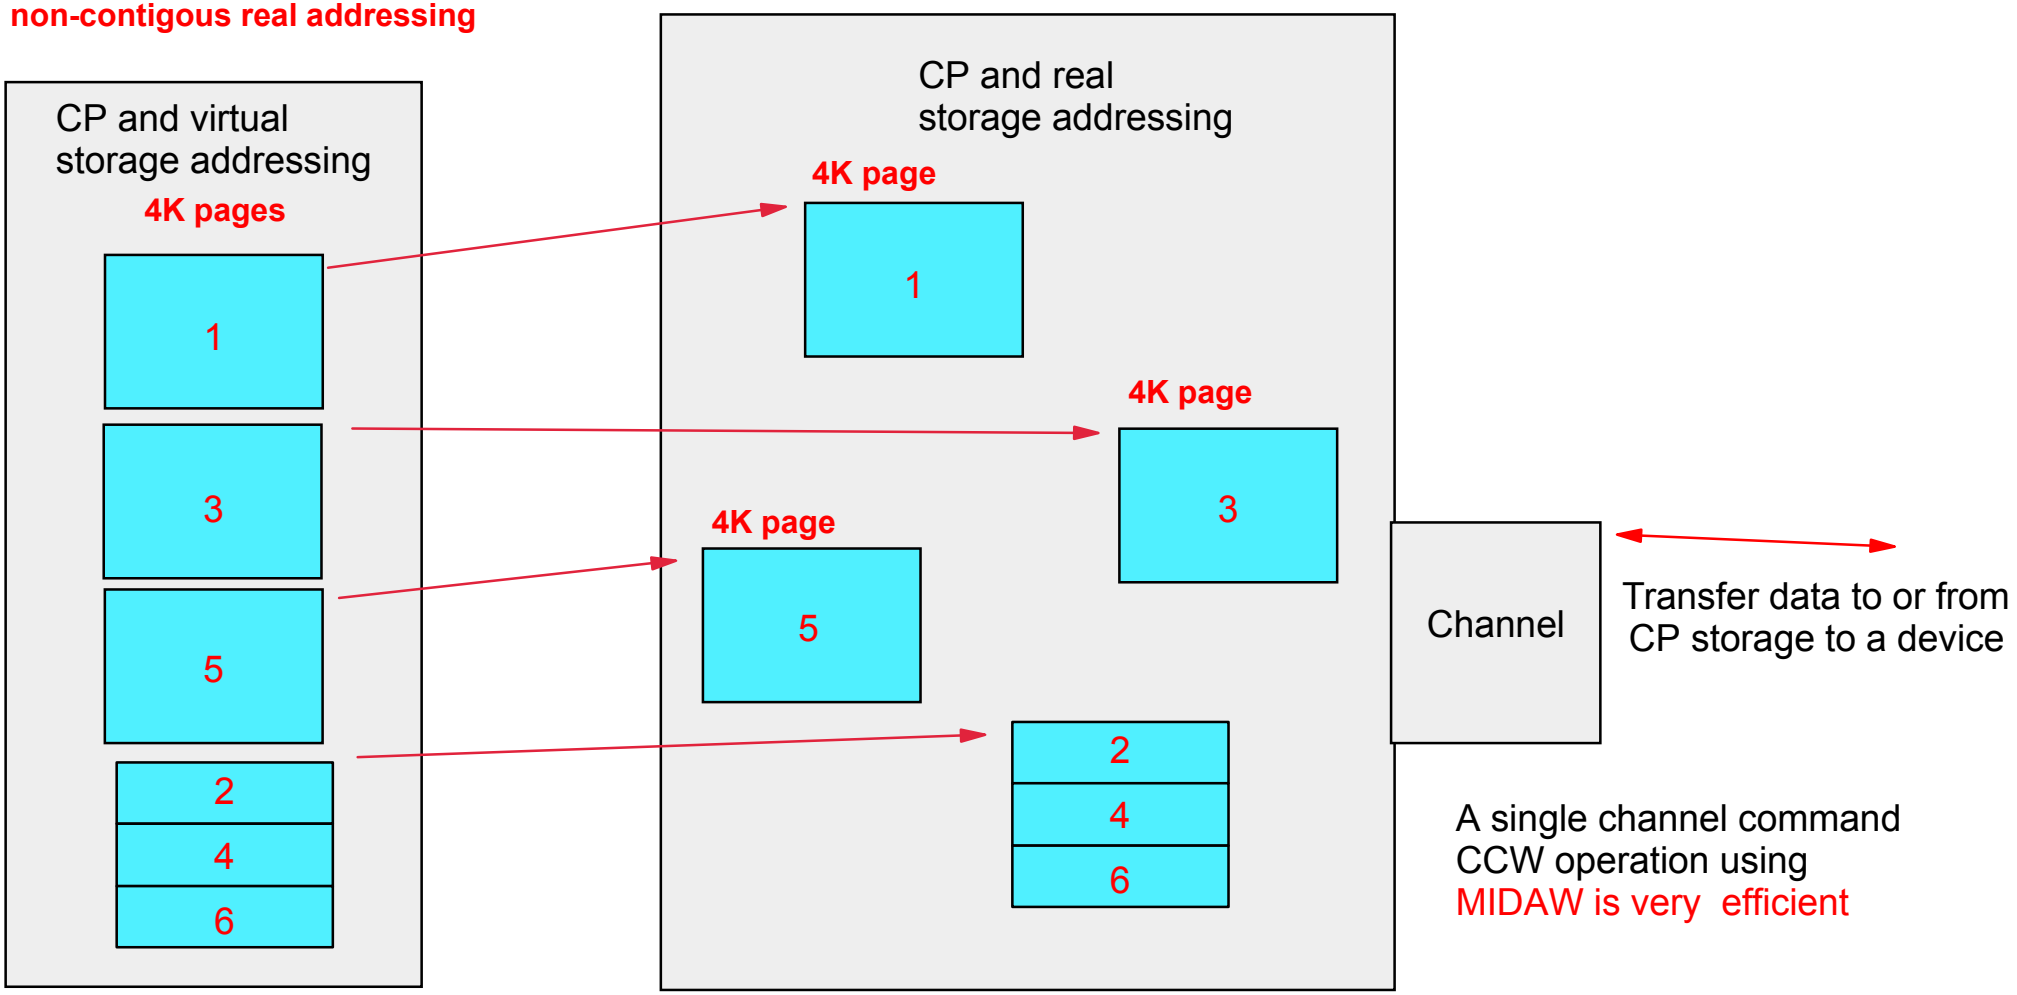

Start Subchannel I/O operation - CCW

All SSCH I/O Operations have a CCW chains. The CCW chain consists of one or more CCWs.

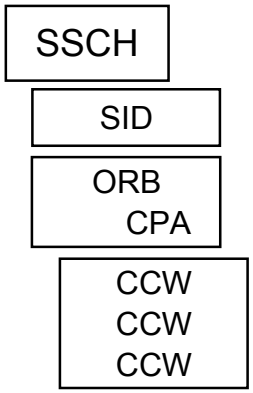

The CCW (Channel Command Word) always has a:

- Command field
- Flags field
- Count field
- Data address field

There are 2 CCW formats

- Format 0 CCW - supports a 24 bit 'real' data addressing field
- Format 1 CCW - supports a 31 bit 'real' data addressing field
- Depending on the CCW Cmd / Flags, the address field points to, another CCW, Data Area, IDAW or a MIDAW

Format 0 CCW

Command	Address (24)	Flags	Count

Format 1 CCW

Command	Flags	Count	Address (31)

- A format 0 CCW IDAW supports 24 bit addressing, Format 1 CCW IDAW supports 31 addressing, and a Format 2 IDA supports 64 bit data addressing
- A MIDAW supports 64 bit data address

CCW Data Address Field - Operation

The meaning of the CCW data address field is dependent on the CCW command and the CCW Flag field bits

Command byte	x'08'	Flag bits - n/a	x'08' command = TIC (Transfers in Channel)	Next CCW address

CCW address = the address of next CCW

Command byte	Not x'08'	Flag bits - x'xx' - Not - IDA or MIDA	CCW data area address

Command byte	x'nn'	Flag bits - x'80' - CD - Chain Data	Current CCW data area address

CCW address = the data area addressed in the next CCW is chained to the current CCW data area

Command byte	x'nn'	Flag bits - x'80' - IDA - Indirect Data Address	Address of 1st IDAW in IDAW list

The CCW address field points to the first IDAW in the IDAW list

Command byte	x'nn'	Flag bits - x'01' - MIDA - Modified Indirect Data Address	Address of 1st MIDAW in MIDAW list

The CCW address field points to the first MIDAW in the MIDAW list

I/O CCW - Data Address - Operations

CCW Data Address

CCW	Cmd	Flag	Count	Data Address

CD (Chain Data)

CCW	Cmd	CD Flag	**CCW Count**	**Data Address**
CCW	--	CD Flag	**CD CCW Count**	**CD Data Address**
CCW		CD Flag	**CD CCW Count**	**CD Data Address**
CCW	--	xx	**CD CCW Count**	**CD Data Address**

IDA operation

CCW	Cmd	IDA Flag	CCW Count	1st IDAW Address

IDAW	IDAW data address Start data transfer within the page	End data transfer at page Boundary
IDAW	IDAW data address Start data transfer at beginning of page	End data transfer at page Boundary
IDAW	IDAW data address Start data transfer at beginning of page	End data transfer at page Boundary
IDAW	IDAW data address Start data transfer at beginning of page	End data transfer within the page at CCW count

MIDA operation

CCW	Cmd	MIDA Flag	CCW Count	1st MIDAW Address

MIDAW	MIDAW Flags	MIDAW Data count	MIDAW Data Address Start data transfer anywhere within the page	End data transfer at MIDAW count or at page Boundary
MIDAW	MIDAW Flags	MIDAW Data Count	MIDAW Data Address Start data transfer anywhere within the page	End data transfer at MIDAW count or at page Boundary
MIDAW	MIDAW Flags	MIDAW Data Count	MIDAW Data Address Start data transfer anywhere within the page	End data transfer at MIDAW count or at page Boundary
MIDAW	MIDAW Flags	MIDAW Dada Count	MIDAW Data Address Start data transfer anywhere within the page	End data transfer at MIDAW count or at page Boundary (or CCW count)

I/O CCW Data Transfer Operation - Non - CD - IDA - MIDAW

A *NON,* CD - IDA - MIDAW single channel-command word (CCW) controls the transfer of data (up to 64K) that spans CONTIGUOUS 2K-byte or 4K-byte 'real' addressing blocks in main storage. For format 0 CCWs the data area real address is from 0M to 16 M and for format 1 CCW the data area real address is from 0M to 2B. Generally applications do not get contiguous real storage greater than 4k in length.

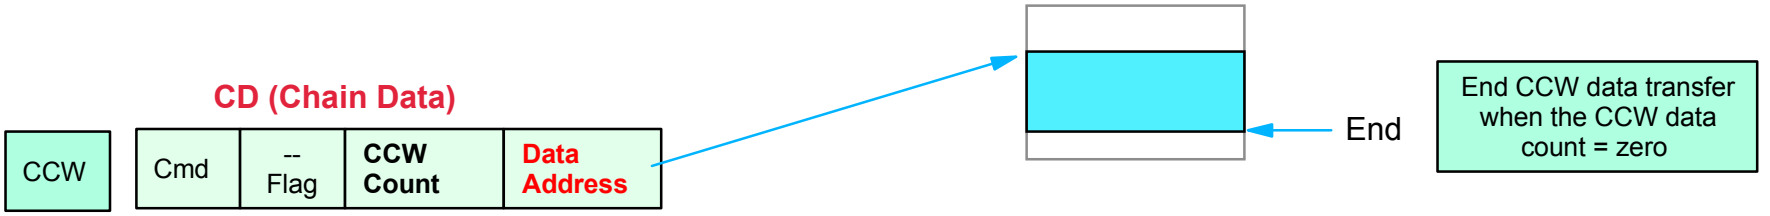

I/O CCW Data Transfer Operation - CD

CCW indirect data addressing (IDA) permits a single channel-command word (CCW) to control the transfer of data that spans noncontiguous 2K-byte or 4K-byte blocks in main storage. The use of CCW indirect data addressing also allows the program to designate data addresses above 16M bytes when using format-0 CCWs or above 2G bytes when using format-1 CCWs.

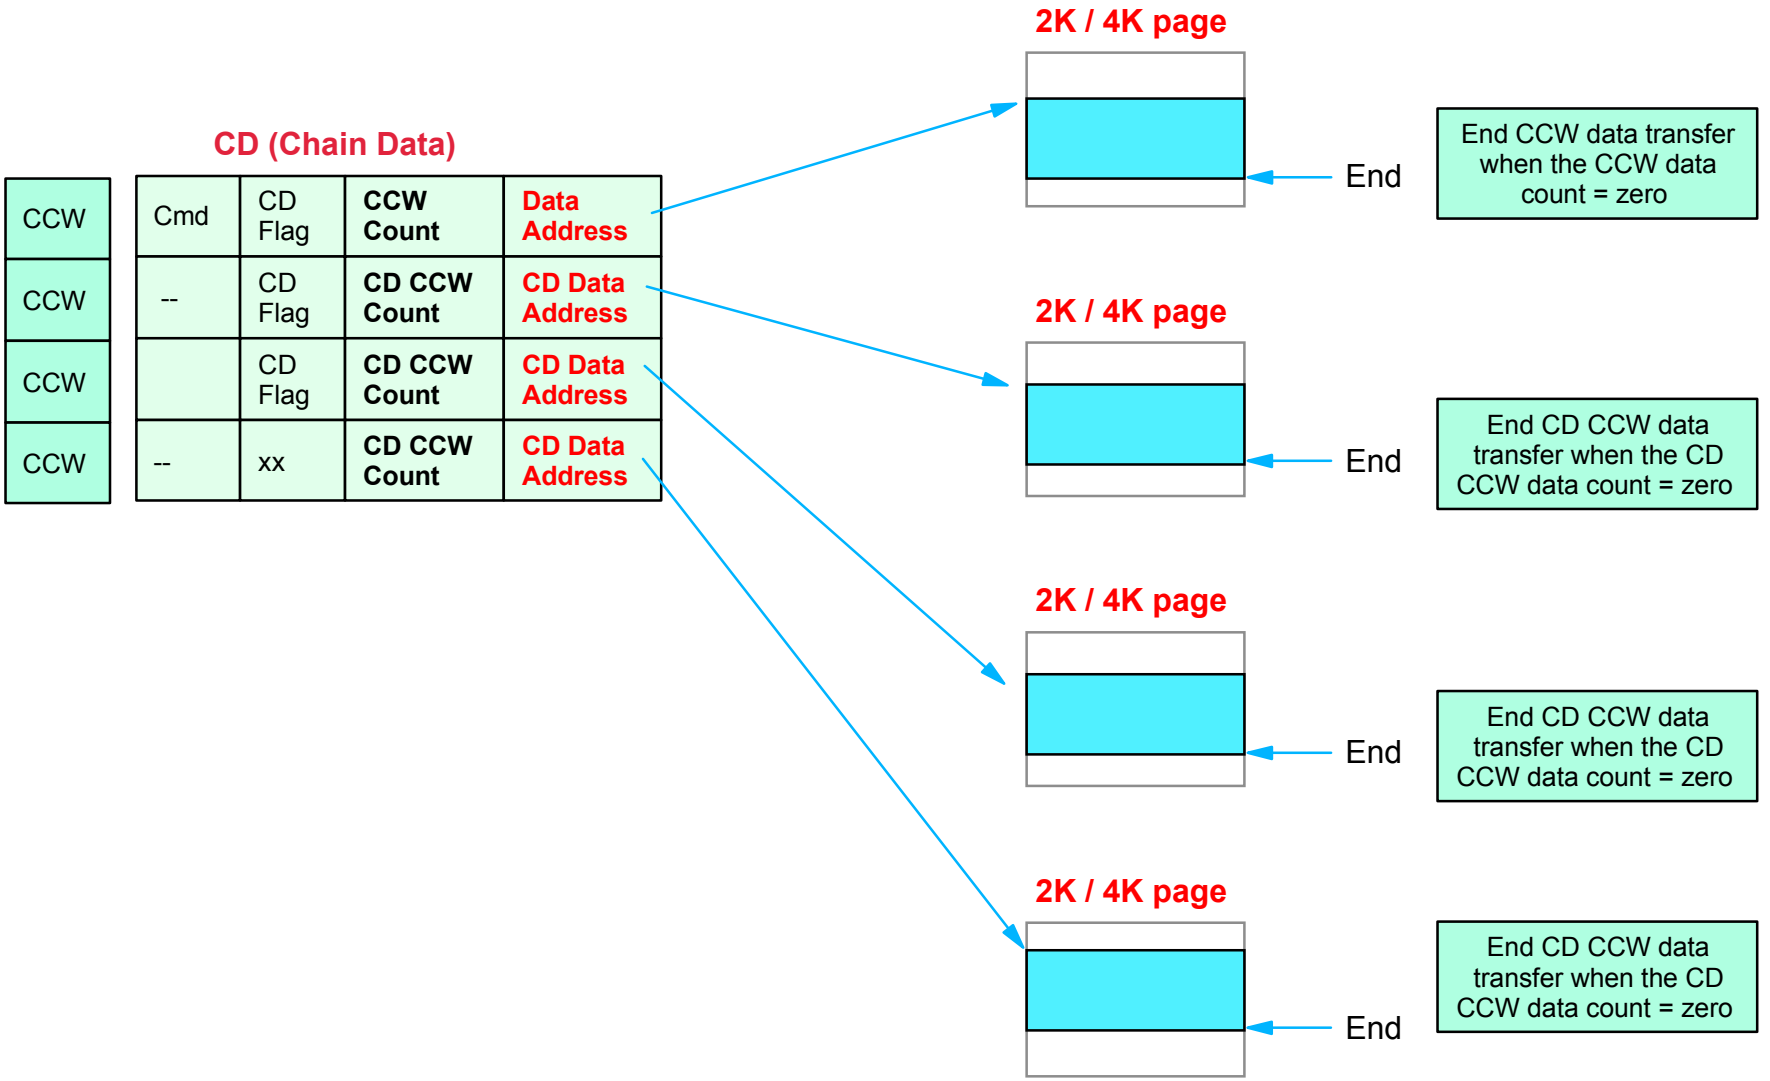

I/O CCW Data Transfer Operation - IDA

CCW indirect data addressing (IDA) permits a single channel-command word (CCW) to control the transfer of data that spans noncontiguous 2K-byte or 4K-byte blocks in main storage. The use of CCW indirect data addressing also allows the program to designate data addresses above 16M bytes when using format-0 CCWs or above 2G bytes when using format-1 CCWs.

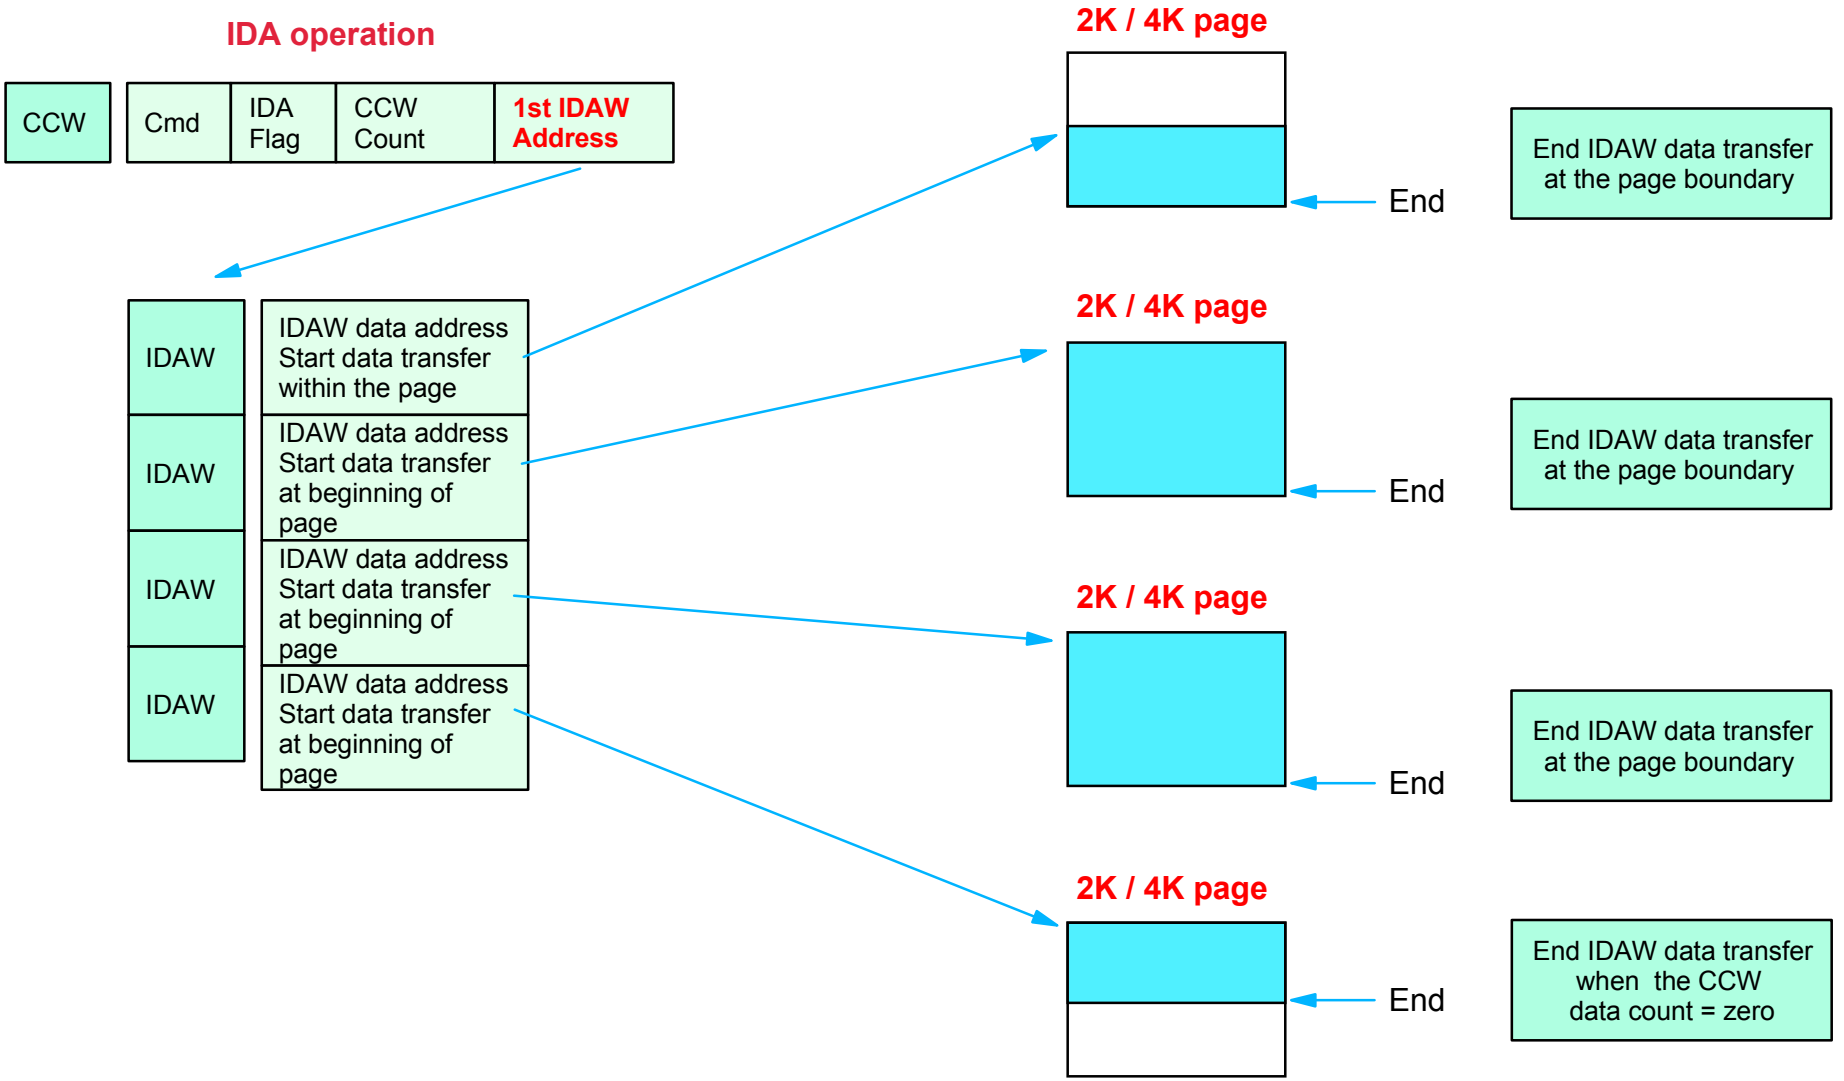

I/O CCW Data Transfer Operations - MIDA

CCW modified indirect data addressing (MIDA) permits a single channel-command word to control the transfer of up to 65,535 bytes of data that spans noncontiguous blocks in main storage. Each block of main storage to be transferred may be specified on any boundary and length up to 4K, provided the specified block does not cross a 4K-byte boundary. The use of modified CCW indirect data addressing requires that the program designate 64-bit data addresses.

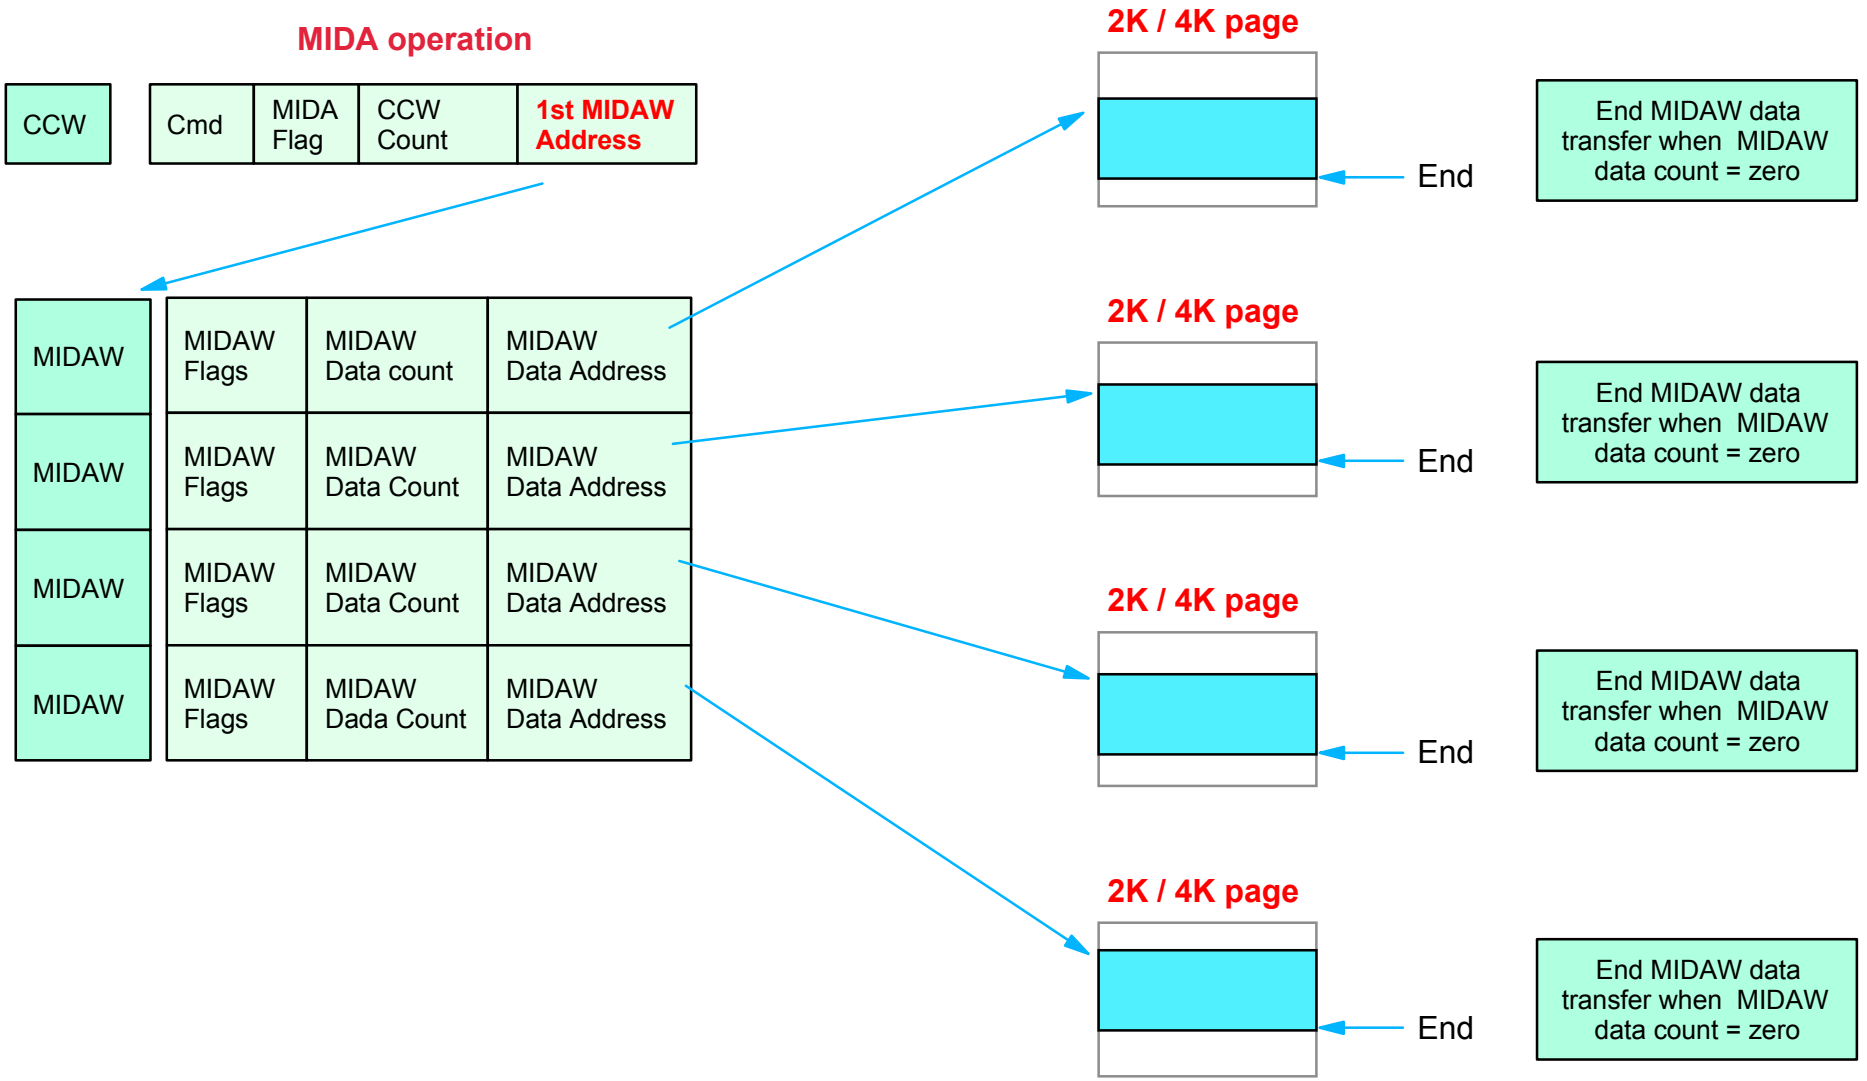

MIDA Facility - MIDAWs

Situation

- z/OS is encouraging the use of CKD larger volumes and data sets
- Extended format data sets provide greater than 4GB addressability
  - Up to now access to extended format datasets uses CCW Chain Data operations
- But extended format data sets today using Data Chaining have a performance penalty
  - Performance penalty is in the channel, switch fabric, and control unit to process the additional CCWs, channel interface sequences, and serial frames when data chaining is used
  - Experiments showed xx% response time degradation and xx% bandwidth reduction under moderate load
- IDAWs cannot be used for Extended Format datasets because of restrictive data addresses
  - The data pointed to by the first IDAW may start anywhere within a 2K of 4K area of storage, but:
  - Subsequent IDAW data areas must start and end on a 2K or 4K boundary (except last)
- MIDAWs provides an alternative to Chain Data operations and eliminates the hardware performance penalties associated with Data Chaining

Example: Data Chaining (CD) CCW chain

```
CCW #1:   Prefix (Write)   (Cmd)   CC   Count = 64
CCW #2:   Read Track Data (Cmd)    CD   Count = 4096   Addr = 20000000
CCW #3:      read track data       CD   Count = 32     Addr = 1E000000
CCW #4:      read track data       CD   Count = 4096   Addr = 25000000
CCW #5:      read track data       CD   Count = 32     Addr = 1E000020
                  .
                  .
                  .
CCW #28:     read track data       CD   Count = 4096   Addr = 31000000
CCW #29:     read track data       --   Count = 32     Addr = 1E0001A0
```

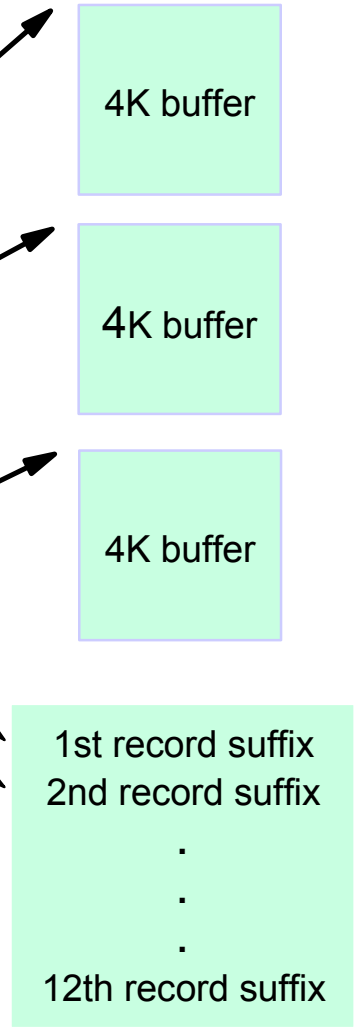

CC = Command Chaining flag
CD = Data Chaining flag

This example of a channel program (non-IDAW non-MIDAW) is used to Read a track's worth of data.
Each record (software record) consists of 4K of data plus a 32 byte suffix.
IDAWs cannot be used since the suffix information does not meet the boundary and length requirements for IDAWs
Each CCW Command Plus Data Chaining CCW command will be send by the channel as a separate (command) sequence to the CU. The sending of all these commands / sequences creates overheads in the channel, fabric and CU. The use of MIDA is eliminate to the sending of the additional command / sequences and reduce he additional overheads in the channel, fabric and CU.

MIDA Facility - MIDAWs

Solution

- Modified Indirect Addressing Facility - referred to as MIDAWs
  - Supported in z/OS 1.6 and 1.7 and on the new z9 2094 processor
    - Availability September 2005
  - MIDAWs channel operations operate like IDAWs
    - A single hardware command / data transfer sequence for the complete CCW
    - Except the boundary and length requirements of IDAW are relaxed for MIDAW
  - A new MIDA control bit (word 1 bit 25 - the 'D' bit) must be on in the ORB
  - New MIDA flag bit (bit 7) must be turned on in CCW flags
    - CCW data address points to a Modified Indirect Address List (MIDAL)
    - Can be used for format 0 or format 1 CCWs (ESAME mode only)
  - Each MIDAL consists of 1 or more MIDAW s - each MIDAW is a quadword
  - Each MIDAW contains the length and address of the data to be used in the I/O operation (CCW operation)
    - Data address can start anywhere within a 4K page
    - Data length can be one byte to 4K bytes -
    - Data length is no longer implied (but must not cross a 4k boundary)

MIDAW

- MIDAW Format
  - Each MIDAW is 16 bytes in length and quadword aligned

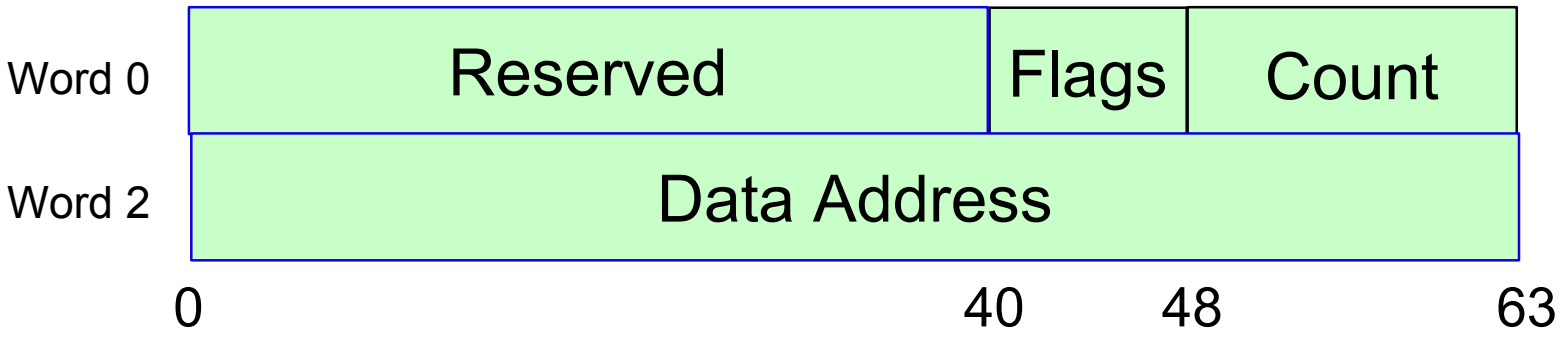

Flags

Bit 0 (40) - Last MIDAW in list

Bit 1 (41) - Skip transfer to main storage (like CCW skip bit)

MIDAW

- Notes:
  - Data length contained in each MIDAW - not implied from storage boundary and length
    - Sum of MIDAW data counts must equal the CCW data count
    - A maximum of 4K can be moved by any one MIDAW - can't cross a 4K page boundary
  - MIDAL must not cross a 4K page boundary (same as IDAL)
  - You can have a mix of IDAWs and MIDAWs in same channel program but not in the same CCW
  - CCW Skip and IDA bits must not be set if MIDA bit is set and vice versa
  - Read skipping is allowed on a MIDAW basis, not a CCW basis
    - Uses MIDAW skip flag
    - When MIDAW read skipping is on, the MIDAW data address is not used but the MIDAW count is included in sum of total MIDAW counts
    - MIDAW Count may exceed 4K for MIDAW with skip flag

SSCH
ORB
MIDA Control
Channel Program Address

CCW
Command
MIDA Flag
MIDAW Total Data Byte Count
MIDAL Address

MIDAW
MIDAW Count
MIDAW Data Address
MIDAW Flags
Data.....

MIDAW
MIDAW Count
MIDAW Data Address
MIDAW Flags
Data.....

MIDAW
MIDAW Count
MIDAW Data Address
MIDAW Flags (Last MIDAW)
Data.....

For writes the data sent by the channel to the CU as one block of data.

For Read the data is received by the channel from the CU as one block of data.

For FICON this means all the data is associated with the same Command IU

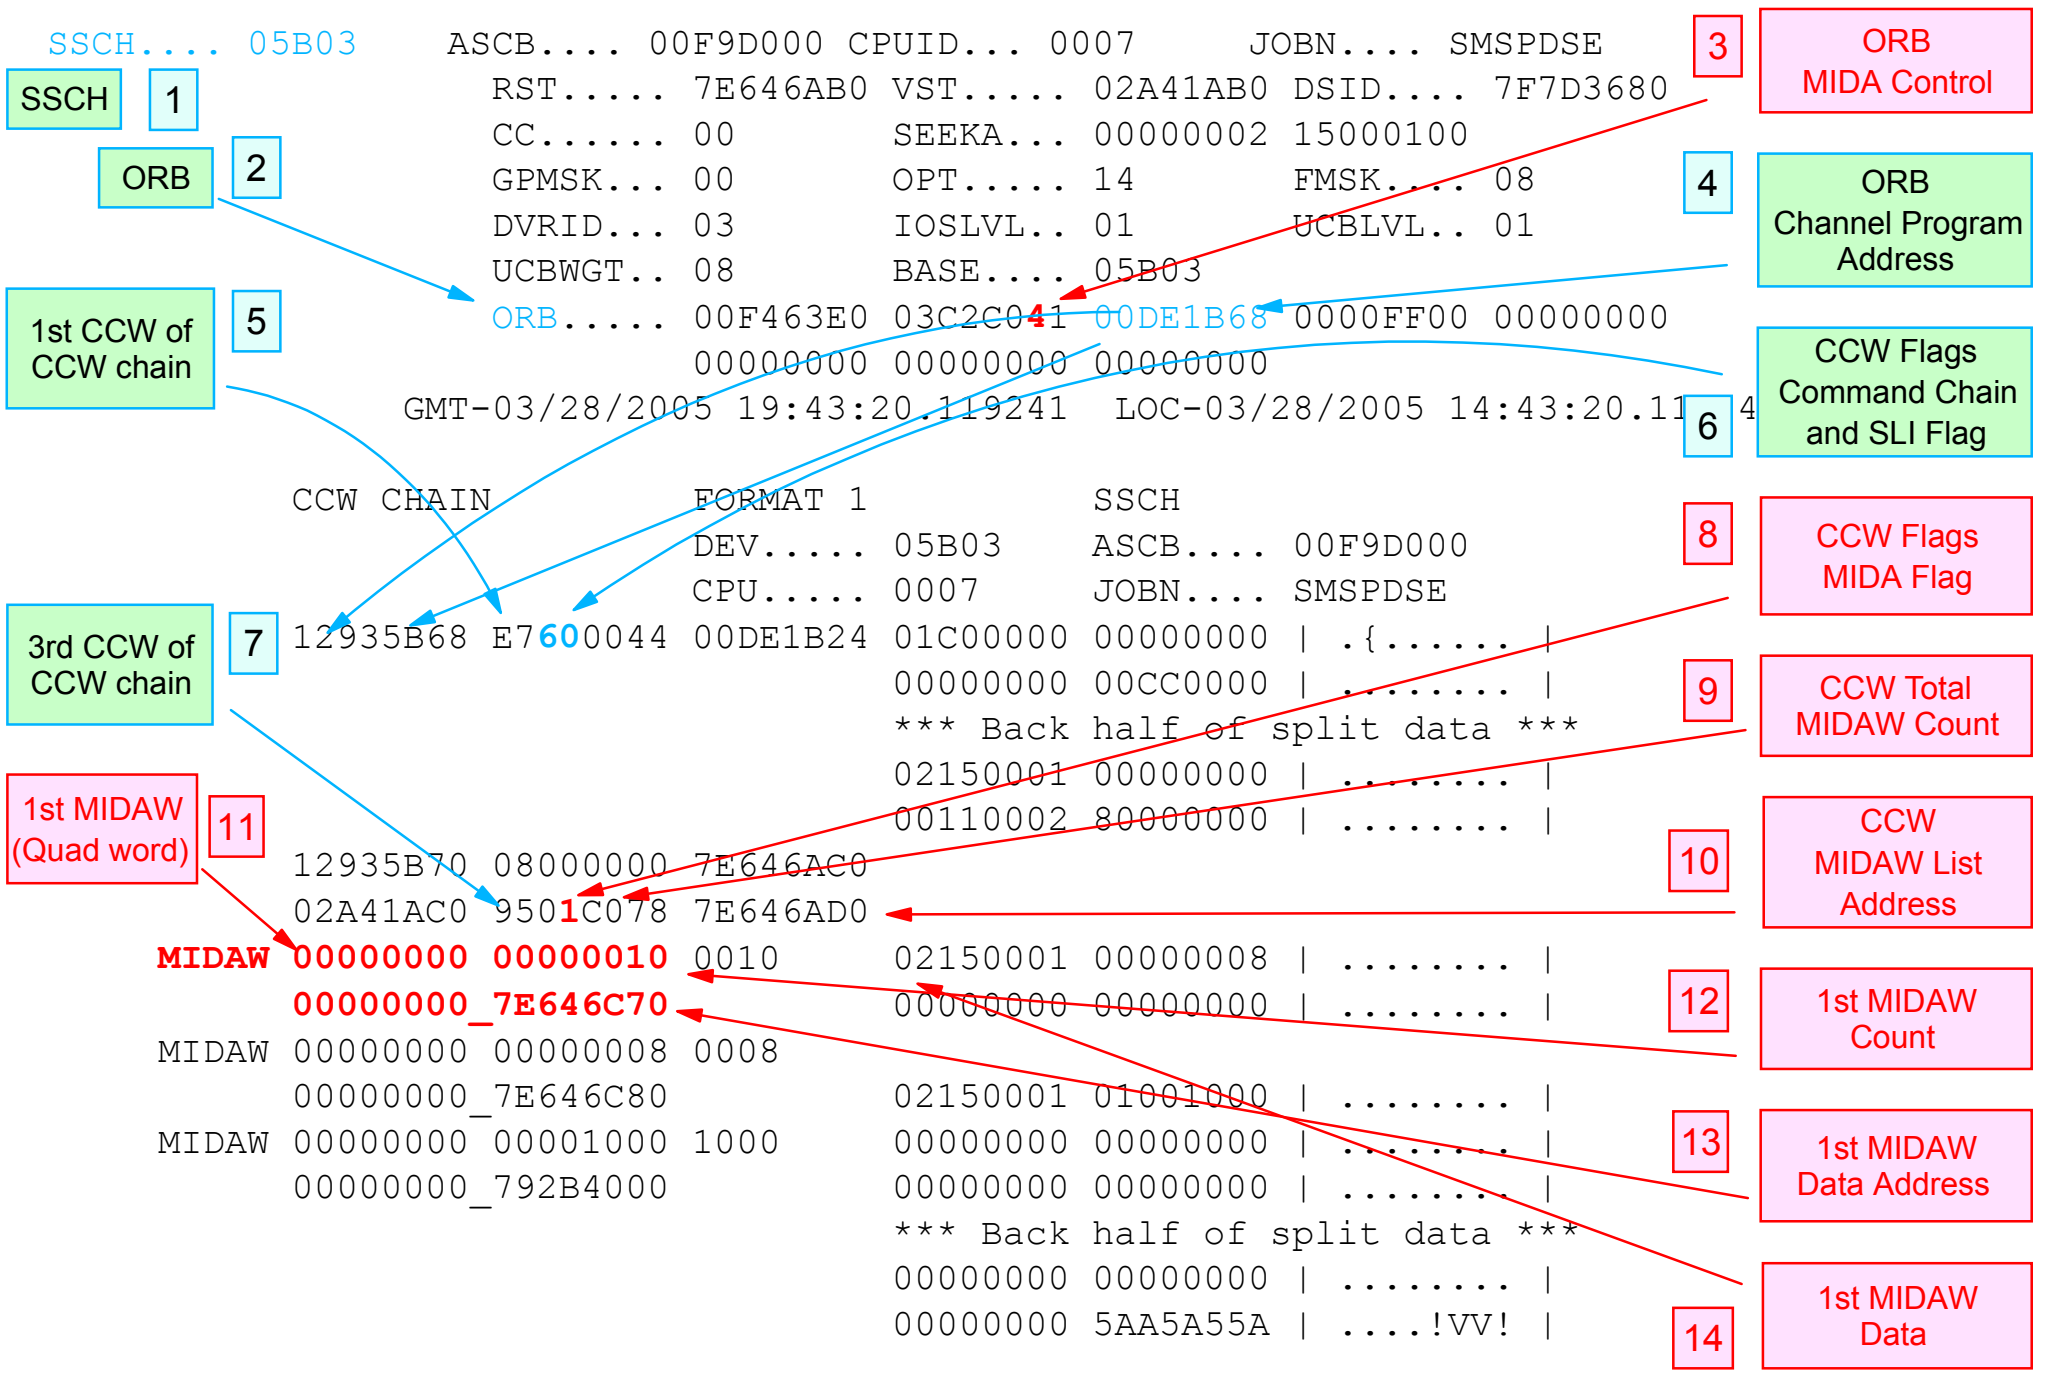
SSCH.... 05B03 ASCB.... 00F9D000 CPUID... 0007 JOBN.... SMSPDSE
RST..... 7E646AB0 VST..... 02A41AB0 DSID.... 7F7D3680
CC..... 00 SEEKA... 00000002 15000100
GPMSK... 00 OPT..... 14 FMSK.... 08
DVRID... 03 IOSLVL.. 01 UCBLVL.. 01
UCBWGT.. 08 BASE.... 05B03
ORB..... 00F463E0 03C2C041 00DE1B68 0000FF00 00000000
00000000 00000000 00000000
GMT-03/28/2005 19:43:20.119241 LOC-03/28/2005 14:43:20.11 4
CCW CHAIN FORMAT 1 SSCH
DEV..... 05B03 ASCB.... 00F9D000
CPU..... 0007 JOBN.... SMSPDSE
12935B68 E7600044 00DE1B24 01C00000 00000000 | .{..... |
00000000 00CC0000 | ..... |
*** Back half of split data ***
02150001 00000000 | ..... |
00110002 80000000 | ..... |
12935B70 08000000 7E646AC0
02A41AC0 9501C078 7E646AD0
MIDAW 00000000 00000010 0010 02150001 00000008 | ..... |
00000000_7E646C70 00000000 00000000 | ..... |
MIDAW 00000000 00000008 0008
00000000_7E646C80 02150001 01001000 | ..... |
MIDAW 00000000 00001000 1000 00000000 00000000 | ..... |
00000000_792B4000 00000000 00000000 | ..... |
*** Back half of split data ***
00000000 00000000 | ..... |
00000000 5AA5A55A | ....!VV! |
SSCH
1
ORB
2
1st CCW of CCW chain
5
3rd CCW of CCW chain
7
1st MIDAW (Quad word)
11
3
ORB MIDA Control
4
ORB Channel Program Address
6
CCW Flags Command Chain and SLI Flag
8
CCW Flags MIDA Flag
9
CCW Total MIDAW Count
10
CCW MIDAW List Address
12
1st MIDAW Count
13
1st MIDAW Data Address
14
1st MIDAW Data

```
SSCH.... 05B03   ORB..... 00F463E0 03C2C041 00DE1B68 0000FF00


CCW  02A41AC0 9501C078 7E646AD0

     MIDAW 00000000 00000010 0010      02150001 00000008 | ........ |
           00000000_7E646C70           00000000 00000000 | ........ |
     MIDAW 00000000 00000008 0008
           00000000_7E646C80           02150001 01001000 | ........ |
     MIDAW 00000000 00001000 1000      00000000 00000000 | ........ |
           00000000_792B4000           00000000 00000000 | ........ |
     MIDAW 00000000 00000008 0008
           00000000_7E646C88           02150001 02001000 | ........ |
     MIDAW 00000000 00001000 1000      00000000 00000000 | ........ |
           00000000_792B4000           00000000 00000000 | ........ |
     MIDAW 00000000 00000008 0008
           00000000_7E646C90           02150001 03001000 | ........ |
     MIDAW 00000000 00001000 1000      00000000 00000000 | ........ |
           00000000_792B4000           00000000 00000000 | ........ |
     MIDAW 00000000 00000008 0008
           00000000_7E646C98           02150001 04001000 | ........ |
     MIDAW 00000000 00001000 1000      00000000 00000000 | ........ |
           00000000_792B4000           00000000 00000000 | ........ |
     MIDAW 00000000 00000008 0008
           00000000_7E646CA0           02150001 05001000 | ........ |
     MIDAW 00000000 00001000 1000      00000000 00000000 | ........ |
           00000000_792B4000           00000000 00000000 | ........ |
     MIDAW 00000000 00000008 0008
           00000000_7E646CA8           02150001 06001000 | ........ |
     MIDAW 00000000 00001000 1000      00000000 00000000 | ........ |
           00000000_792B4000           00000000 00000000 | ........ |
```

- ORB MIDA control = 1
- CCW MIDA flag = 1
- The total MIDAW data transfer count should be equal to the CCW data count when there is no SILI flag
- Add all the MIDAW counts on foil 1/2 and 2/2
- For this example:

  The CCW data count count is x'C078'

  The total MIDAW data count is x'C078'

  The complete data will be sent as one block of data from the channel to the CU (49272 bytes)

  For FICON this will be 7 IUs (4 information units)

Foil 2 of 2

```
MIDAW 00000000 00000008 0008
      00000000_7E646CB0          02150001 07001000 | ........ |
MIDAW 00000000 00001000 1000     00000000 00000000 | ........ |
      00000000_792B4000          00000000 00000000 | ........ |
MIDAW 00000000 00000008 0008
      00000000_7E646CB8          02150001 08001000 | ........ |
MIDAW 00000000 00001000 1000     00000000 00000000 | ........ |
      00000000_792B4000          00000000 00000000 | ........ |
MIDAW 00000000 00000008 0008
      00000000_7E646CC0          02150001 09001000 | ........ |
MIDAW 00000000 00001000 1000     00000000 00000000 | ........ |
      00000000_792B4000          00000000 00000000 | ........ |
MIDAW 00000000 00000008 0008
      00000000_7E646CC8          02150001 0A001000 | ........ |
MIDAW 00000000 00001000 1000     00000000 00000000 | ........ |
      00000000_792B4000          00000000 00000000 | ........ |
MIDAW 00000000 00000008 0008
      00000000_7E646CD0          02150001 0B001000 | ........ |
MIDAW 00000000 00001000 1000     00000000 00000000 | ........ |
      00000000_792B4000          00000000 00000000 | ........ |
MIDAW 00000000 00000008 0008
      00000000_7E646CD8          02150001 0C001000 | ........ |
MIDAW 00000000 00001000 1000     00000000 00000000 | ........ |
      00000000_792B4000          00000000 00000000 | ........ |
                                 *** Back half of split data ***
                                 00000000 00000000 | ........ |
                                 00000000 5AA5A55A | ....!VV! |
      00000000 00800008 0008
      00000000_7E646CE0          FFFFFFFF FFFFFFFF | ........ |
IO...... 05B03    ASCB.... 00F9D000 CPUID... 0005     JOBN.... SMSPDSE
                  PSW..... 070E0000 00000000
                  IRB..... 00C04007 7E646AC8 0C000000 0040002F 00000000
```

Last MIDAW in MIDAW list

The total MIDAW data count should be equal to the CCW data count as there is no SILI flag..

For this example:

The CCW data count count is x'C078'

The total MIDAW data count is x'C078'

Data transfer is complete when (which ever comes first) the MIDAW data tranfer count is equal to the CCW command count or after processing of a MIDAW that has the 'last MIDAW flag bit' set

I/O interrupt shows that the I/O operation completed successfully with a status of CE/DE (x'0C') and a residual count of x'0000'

The addition of all the MIDAW data counts was equal to CCW count in the MIDA flag CCW

CD and IDA, and CD and MIDA

- A CD condition and IDA flag may be be *in operation* for the same CCW
  - A FICON Cmd Data IU (Write) or Cmd IU (Read) will be sent for the CD CCW
    - CCW # 1. Normal ComandData IU + CD flag = Normal data transfer
    - CCW # 2. CD CCW = CD command, if IDA = 1 then IDA data transfer
    - Continue IDA data transfer until count = zero, or no more data to send / receive

- A CD condition and MIDA flag may be be *in operation* for the same CCW
  - A FICON Cmd Data IU (Write) or Cmd IU (Read) will be sent for the CD CCW
    - CCW # 1. Normal ComandData IU + CD flag = Normal data transfer
    - CCW # 2. CD CCW = CD command, if MIDA = 1 then MIDA data transfer
    - Continue MIDA data transfer until count = zero, or no more data to send / receive

z/OS MVS Commands for MIDAW

- D IOS,MIDAW
  - Displays the current system support status (z/OS and z9 or zSeries) of MIDAW
- SETIOS MIDAW=NO
  - Turns off z/OS MIDAW support
- SETIOS MIDAW=YES
  - Turns on z/OS MIDAW support
  - Requires z9 Server support
- DS QDASD,dddd,UCB
  - Display the DASD UCB and the UCB MIDAW support bit

z/OS MVS Commands for MIDAW (on a z9 109)

```
D IOS,MIDAW
IOS097I 11.47.04 MIDAW FACILITY 790
MIDAW FACILITY IS ENABLED


SETIOS MIDAW=NO
IOS090I SETIOS. MIDAW UPDATE(S) COMPLETE


D IOS,MIDAW
IOS097I 11.50.13 MIDAW FACILITY 794
MIDAW FACILITY IS DISABLED


SETIOS MIDAW=YES
IOS090I SETIOS. MIDAW UPDATE(S) COMPLETE


D IOS,MIDAW
IOS097I 11.53.13 MIDAW FACILITY 801
MIDAW FACILITY IS ENABLED
```

z/OS MVS Commands for MIDAW (on a z9 109)

```
D M=DEV(dddd)
```

z9 and z/OS MIDAW enabling

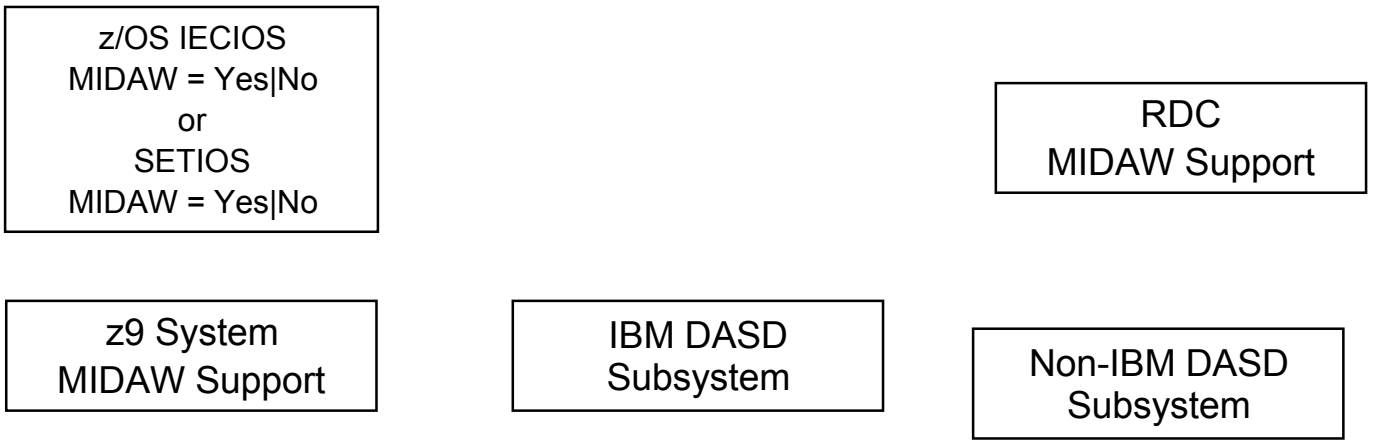

z/OS MVS Commands for MIDAW (on a z9 109)

```
D IOS,MIDAW
IOS097I 12.03.44 MIDAW FACILITY 813
MIDAW FACILITY IS ENABLED


DS QDASD,8103,UCB
IEE459I 12.03.49 DEVSERV QDASD 815
UNIT VOLSER SCUTYPE DEVTYPE   CYL  SSID SCU-SERIAL DEV-SERIAL EF-CHK
8103 BH7CD1 2105800 2105000  3339  8911 0113-22513 0113-22513 **OK**
  UCB AT V00F201A0
0088FF8481030000 0000000008E4C3C2 3030200F00F20178 002D0100C2C8F7C3
C4F1100000A00000 00F1FFA0021715A8 0380030300000000
  UCB PREFIX AT V02417330
000C804000000000 00000000000107DB 289C1FE4C10040FF 8081848582838687
0108000000000001
  UCB COMMON EXTENSION AT V00F20178
00000940202A0008 0241733000000000 0000000000FCEC0C 00F2013800001A00
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

z/OS MVS Commands for MIDAW (on a z9 109)

```
SYS1.PARMLIB(IECIOSxx)

MIDAW=NO|YES
```

```
T IOS=xx

D IOS,MIDAW
IOS097I 12.03.44 MIDAW FACILITY 813
MIDAW FACILITY IS ENABLED
```

End of MIDAW

Presentation