

ADABAS and Disk Usage

Is it time to reevaluate?

Becky Albin
IT Architect
Software AG

Becky.Albin@softwareagusa.com

Agenda

- 'Large' datasets
 - IBM requirements
 - Adabas use
- Number of ASSO/DATA datasets
 - PLOG/CLOG
- Block Sizes
- Performance / Impact / Considerations

Large Datasets

Adabas and 'Large' Datasets

- IBM now allows more than 65k tracks per dataset
 - z/OS 1.7
 - Adabas v7.4.4
 - Release Notes, page 18
- Adabas datasets can now occupy an entire volume
- Datasets are still limited to 59 volumes
- Datasets can have up to 16 extents on a volume

Adabas and 'Large' Datasets

- PARMLIB parameter: BLOCKTOKENSIZE(REQUIRE) *
- JCL DD statements used to allocate the 'large' datasets must specify DSNTYPE=LARGE *
 - ADAFRM required
 - No longer needed after allocation on other DD statements
- Applies to:
 - ASSO/DATA/WORK
 - TEMP/SORT
 - PLOG/CLOG
 - RLOG
 - DSIM
- Use ADASAV SAVE / RESTORE or ADAORD RESTRUCTUREDB / STORE from original datasets to new larger datasets

* *be sure that any offsite/remote facility supports this feature*

Adabas and Large Blocks on 'Tape'

- Exploits IBM's large block (>32,760 bytes) support for sequential access methods BDAM and QSAM
 - OS/390 v2 r10 and an z/OS release
 - ADAIOR supports tape drives with a block sizes of up to:
 - 256k for 3590 devices
 - 64k for 3490/3490E devices
- Provides performance benefits for any Adabas utility writing to tape, such as ADASAV
- If shipping tapes to off-site facilities, the off-site center must support these 'large blocks'

Adabas - Increased number of DATA/ASSO datasets

Adabas and Physical Extents/Datasets

- Adabas v8.1 supports up to 99 datasets for:
 - DATA
 - ASSO

```
//DDDATAR1 DD DSN=ADABAS.DB54.DATAR1,DISP=SHR
```

```
//DDDATAR2 DD DSN=ADABAS.DB54.DATAR2,DISP=SHR
```

```
...
```

```
...
```

```
//DDDATAR15 DD DSN=ADABAS.DB54.DATAR15,DISP=SHR
```

Adabas and Physical Extents/Datasets

- Adabas V7.4 introduced up to 8 PLOG/CLOG datasets

```
ADARUN NPLOG=n,PLOGDEV=n,PLOGSIZE=n  
ADARUN NCLOG=n,CLOGDEV=n,CLOGSIZE=n  
UEX12=name
```

or

```
ADARUN DUALPLD=3390,DUALPLS=n  
ADARUN DUALCLD=3390,DUALCLS=n  
UEX2=name
```

```
//DDPLOGR1 DD DSN=ADABAS.DB54.PLOGR1,DISP=SHR  
//DDPLOGR2 DD DSN=ADABAS.DB54.PLOGR2,DISP=SHR  
...  
...  
//DDPLOGR4 DD DSN=ADABAS.DB54.PLOGR4,DISP=SHR
```

Adabas - Block Sizes

Adabas Block Sizes

- Why do we care about block sizes?
 - Bigger vs Smaller
- Performance
 - LBP
 - LFIOP
 - ADACSH
 - LP
 - PLOG
- Bigger records/Spanned records
 - Large Objects
- More Logical Extents
 - DS/AC/UI/NI/DSST
- More Physical Extents
 - ASSO/DATA

Device	Trks Cyl	ASSO	DATA	WORK	PLOG RLOG	CLOG	TEMP SORT DSIM	*
3390	15	2544:18	5064:10	5724:9	5724:9	5064:10	8904:6	3
8390	15	3440:14	6518:8	10706:5	10706:5	8904:6	8904:6	1
8391	15	4136:12	10796:5	13682:4	13682:4	8904:6	18452:3	1
8392	15	4092:12	12796:4	18452:3	18452:3	18452:3	18452:3	1
8393	15	4092:12	27644:2	27990:2	27990:2	27990:2	27990:2	1

1) The 8350, 838*n*, and 839*n* are pseudodevice types physically contained on a 3350, 3380, and 3390 device, respectively, but for which some or all of the standard **≤block sizes≥** are larger.

3) The IBM RAMAC 9394 emulates devices 3390 Model 3, 3380 Model K, or 9345 Model 2.

Adabas Block Sizes

- Why do we care about block sizes?
- Might just want more records per block
- Caution when using different block sizes across multiple datasets on a single database
 - Ex: //DDDATAR1 is allocated as 3390 (5064 bytes)
//DDDATAR2 is allocated as 8391 (10796 bytes)
 - Better to use same block size across datasets?
- Zap ADARUN defaults (MVSJOBS – DEFAULTS)
 - DEVICE=3390 – refers to DDASSOR1
 - GCB

Adabas Block Sizes

- Maximum compressed record length is DS block size - 4 bytes
 - With v8, can be much larger

```

ADADBS RECORDSPANNING FILE = file-number
                        MODE = { ON | OFF }
                        [PASSWORD = 'password' ]
                        [NOUSERABEND]
                        [TIMELIMIT = { 60 | timelimit } ]
                        [TEST]
    
```

- Large Object Fields:
 - Single field up to 2G in size (2,147,483,643 bytes)
 - Binary and Character LB fields
 - Contained in separate Adabas file
 - Could be different block size than 'base' file

Adabas Block Sizes

- Remember that WORK & PLOG must be larger than DS
 - at least 100 bytes
- CLOG must be = or > DS
- Always use Forward Index Compression for ALL files
- Standard vs Pseudo vs User-Defined
 - stick with 1st or 2nd option
 - User-defined can be done – but...
 - Make multiple of 256 bytes
 - make ASSO a multiple of DATA

Adabas Block Sizes

NO SPANNED RECORDS



BH

BLOCK HEADER 4b

2b – INCLUSIVE
USED BLOCK
LENGTH

2b – x'0004'

RH

RECORD HEADER 6b

2b – INCLUSIVE
RECORD
LENGTH

4b – ISN


fields

FIELDS:

LENGTHS
MU/PE COUNTERS
EMPTY FLD COUNTS
VALUES

ONE SPANNED RECORD

DS RABN:



SPANNED RECORD HEADER 20b:

2b – INCLUSIVE RECORD LENGTH

4b – UNKNOWN

4b – ISN OF THIS RECORD

2b – INDICATOR: x'0081' Primary Record

x'0041' Secondary Record

4b – ISN OF NEXT SECONDARY (x'00's IF LAST)

4b – ISN OF PRIMARY RECORD

Adabas v8.1 Spanned Records

1 PRIMARY RECORD + UP TO 4 SECONDARY RECORDS

RECORD SIZE LIMITED BY COMPRESSED RECORD SIZE

(RSP.CD 49) :

PRIOR V8:

**COMPRESSED RECORD
DS BLOCKSIZE:**

3390: 5,064

8391: 10,796

8393: 27,644

V8:

**COMPRESSED RECORD
DS BLOCKSIZE:**

3390: 5,064 X 5 = 25,320

8391: 10,796 X 5 = 53,980

8393: 27,644 X 5 = 138,220

**NOTE: ACTUAL DATA LIMITED BY BLOCK HEADER,
RECORD HEADER, SPAN HEADER**

Adabas Block Sizes

- Impact on LBP and LFIOP:
 - LBP of same size will hold less blocks if using larger block size
 - ex: v74 LBP=50M using 'standard' block sizes will be allocated with ~ 1meg for headers and ~ 49 meg for ASSO/DATA blocks
 - 19,200 ASSO blocks or
 - 9,600 DATA blocks
 - Same LBP=50M using 8391 block sizes will be allocated with ~ 600k for headers and ~ 49 meg for ASSO/DATA blocks
 - 11,800 ASSO blocks or
 - 4,500 DATA blocks
 - Performance can be reduced significantly due to increase in physical I/Os
 - Will need to increase LBP if using larger block sizes
 - Increase LFIOP or buffer flushes will happen more often

Adabas Block Sizes

- Using ADACSH?
 - Read-ahead option for sequential commands (L3, L2, L9, non-de)
 - Reads all consecutive blocks on a single track
 - Brought into CSH area
 - CBUFNO and CEXCLUDE parameters
- FastPath also does read-ahead caching
- LP (Work Part 1) may need to be increased, should be large anyway
 - RC 9 - sc 15
- PLOGSIZE may need to be increased, due to block images written during online backups
 - Or add additional PLOG datasets


Adabas Block Sizes

- Spanned records impact on ADARUN parameters:
 - NH/NISNHQ
 - LWP
 - LP/PLOGSIZE

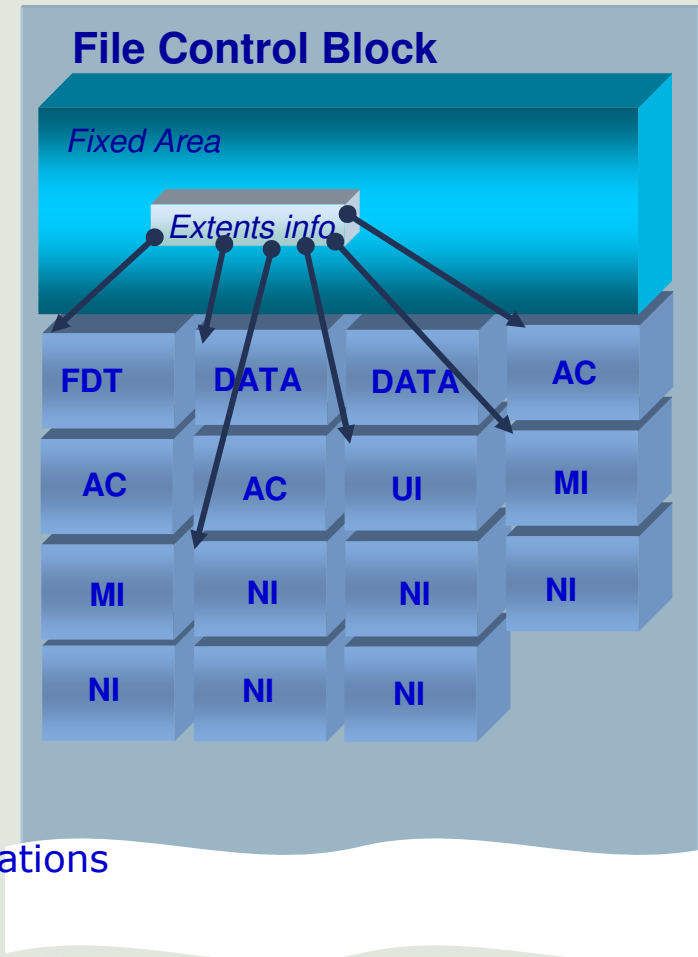
Adabas Block Sizes

- Moving from 'Standard' to 'Pseudo'
 - ADAAORD RESTRUCTURE / STORE
 - backups ???
- while doing this consider:
 - 4 byte RABNs
 - 4 byte ISNs
- ADASAV RESTORE

Adabas Block Sizes & Logical Extents



- one block in Asso
- Fixed area location
- Extent RABNs
- File features
 - MU/PE
 - Spanned DS
 - LOB support
- FDT location
- File extent info
- LOB Info



- Increase logical file extents for NI, UI, DS allocations
- Maximum depends on size of Associator block
 - 3390 can have around 40 extents of each type
 - Or there could be more of one type and less for another
- Removing five-extent limits provides better 24x7 support

Adabas Block Sizes

- Other things to ponder...
 - File placement ???
 - usage
 - volume
 - controller
 - Allocating 'dummy' files