Informix Best Practices Configuration, ONCONFIG, CPU and Memory Usage

Webcast – February 23, 2017

by

Lester Knutsen

Advanced DataTools

Lester Knutsen

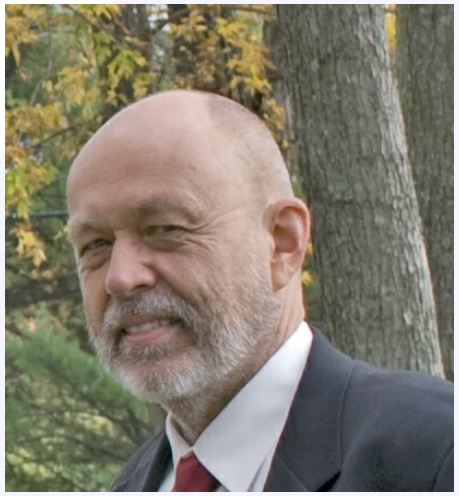

Lester Knutsen is President of Advanced DataTools Corporation, and has been building large Data Warehouse and Business Systems using Informix Database software since 1983. Lester focuses on large database performance tuning, training and consulting. Lester is a member of the IBM Gold Consultant program and was presented with one of the Inaugural IBM Data Champion awards by IBM. Lester was one of the founders of the International Informix Users Group and the Washington Area Informix User Group.

lester@advancedatatools.com
www.advancedatatools.com
703-256-0267 x102

Overview

- CPU Recommendations and Best Practices
- Memory Recommendations and Best Practices
- ONCONFIG Recommendations and Best Practices
  - Basic Settings
  - Additional Key Settings

CPU – Central Processor Unit Recommendations for Informix and Best Practices

Advanced DataTools

CPU Terms

- Socket = One Chip or Processor
- Cores per Socket = How many cores run on a chip. A core only runs one process at a time.
- Hyper-Threads or SMT threads per Core = Many Cores have the ability to run multiple threads. No matter how many threads run on a Core, only one thread can run at a time on a core. Hyper-Threads will appear as additional Virtual Cores.
- Chip speed is measured in gigahertz (GHz); this is the speed of a single core of your processor.
- PVU - IBM Processor Value Unit = A unit of measure used to differentiate licensing of software

CPU Terms

- Example: 2 Sockets with 5 Cores and 2 Hyper-Threads per Core = 10 Cores and 20 Virtual Cores
- Can run 10 processes at the same time

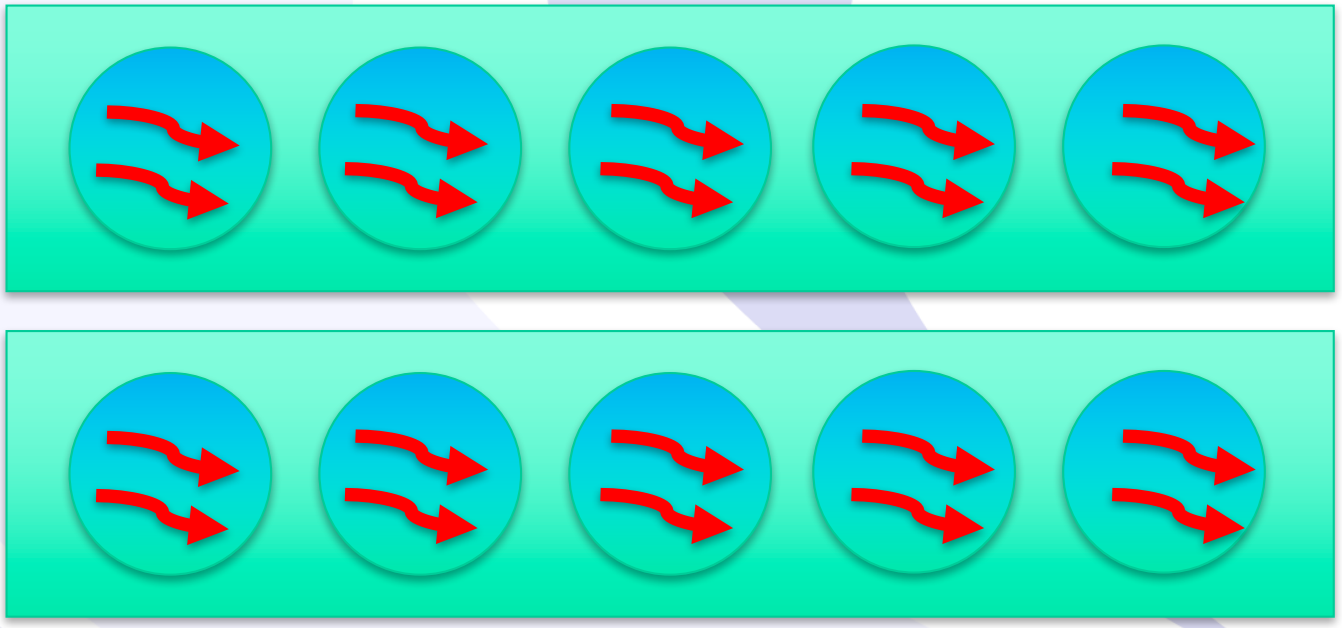

Informix CPU Best Practices

- How many Cores will be allocated for Informix? What else is running on the machine?
- Traditional best practice is number of physical CPU Cores minus 1
- Current CPU Cores are fast enough to handle 2-3 oninits per Core or 1 oninit per 500-1000 MHz

CPU Usage Best Practices

- How busy are your CPUs?
- Tools to monitor:
  - sar –u, vmstat, mpstat, top, prstat
- Performance Guideline for Average CPU Usage:
  - < 30 % - Good
  - 30-60% - Fair
  - > 60% - Poor
- ***Save 60% of your CPU usage to handle Workload Spikes or you may not be able to handle busy loads***

Hyper-Threads or SMT Threads

- Hyper-Threads and SMT Threads may not be helpful to Informix oninit process
- Example:

Informix on AIX. Each AIX Core has 4 SMT threads, each displaying as a CPU; only the first thread is busy, the rest are idle.

```
System configuration: lcpu=256 mode=Capped

cpu  min  maj  mpc  int   cs  ics   rq  mig lpa sysc us sy wa id   pc
  0 2457    0    0  447  173    9    0   90  91 11091  9 34  0 57 0.34
  1    0    0    0   54    0    0    0    0   -    0  0  0  0 100 0.22
  2    0    0    0   54    0    0    0    0   -    0  0  0  0 100 0.22
  3    0    0    0   61    0    0    0    0 100    0  0  3  0 97 0.22
  4  843    0    0  665  670   66    0  539  90 10674 14 16  0 70 0.30
  5    0    0    0   60    3    1    0    4  83    5  0  4  0 96 0.23
  6    0    0    0   59    0    0    0    1 100    0  0  4  0 96 0.23
  7    0    0    0   60    3    0    0    4 100    8  0  3  0 97 0.23
  8  413    0    0  504  695   57    0  533  91 15554 16 20  0 64 0.32
  9    0    0    0   60   15    1    0   14  96   36  0  5  0 95 0.23
 10    0    0    0   95   28    2    0    4 100  145  1  2  0 97 0.23
 11    0    0    0   59    0    0    0    0 100    0  0  4  0 96 0.23
 12  466    0    0  494  648   38    0  451  91 12710 17 12  0 71 0.30
 13    0    0    0   61    0    0    0    1  88    1  0  6  0 94 0.24
 14    0    0    0   60    0    0    0    5 100    1  0  5  0 95 0.24
 15    0    0    0   54    0    0    0    2 100    1  0  0  0 100 0.23
```

Hyper-Threads or SMT Threads Best Practices

- Test, Test and Test again; don't assume that more Hyper-Threads or SMT threads are better. Your workload will determine what is best.
- AIX – Try 2 SMT threads per Core on Power6 and Power7, 4 SMT on Power8
- Intel – Try 2 Oninits per Core instead of 2 Hyper-Threads and 1 Oninit per Hyper-Thread

Informix Architecture

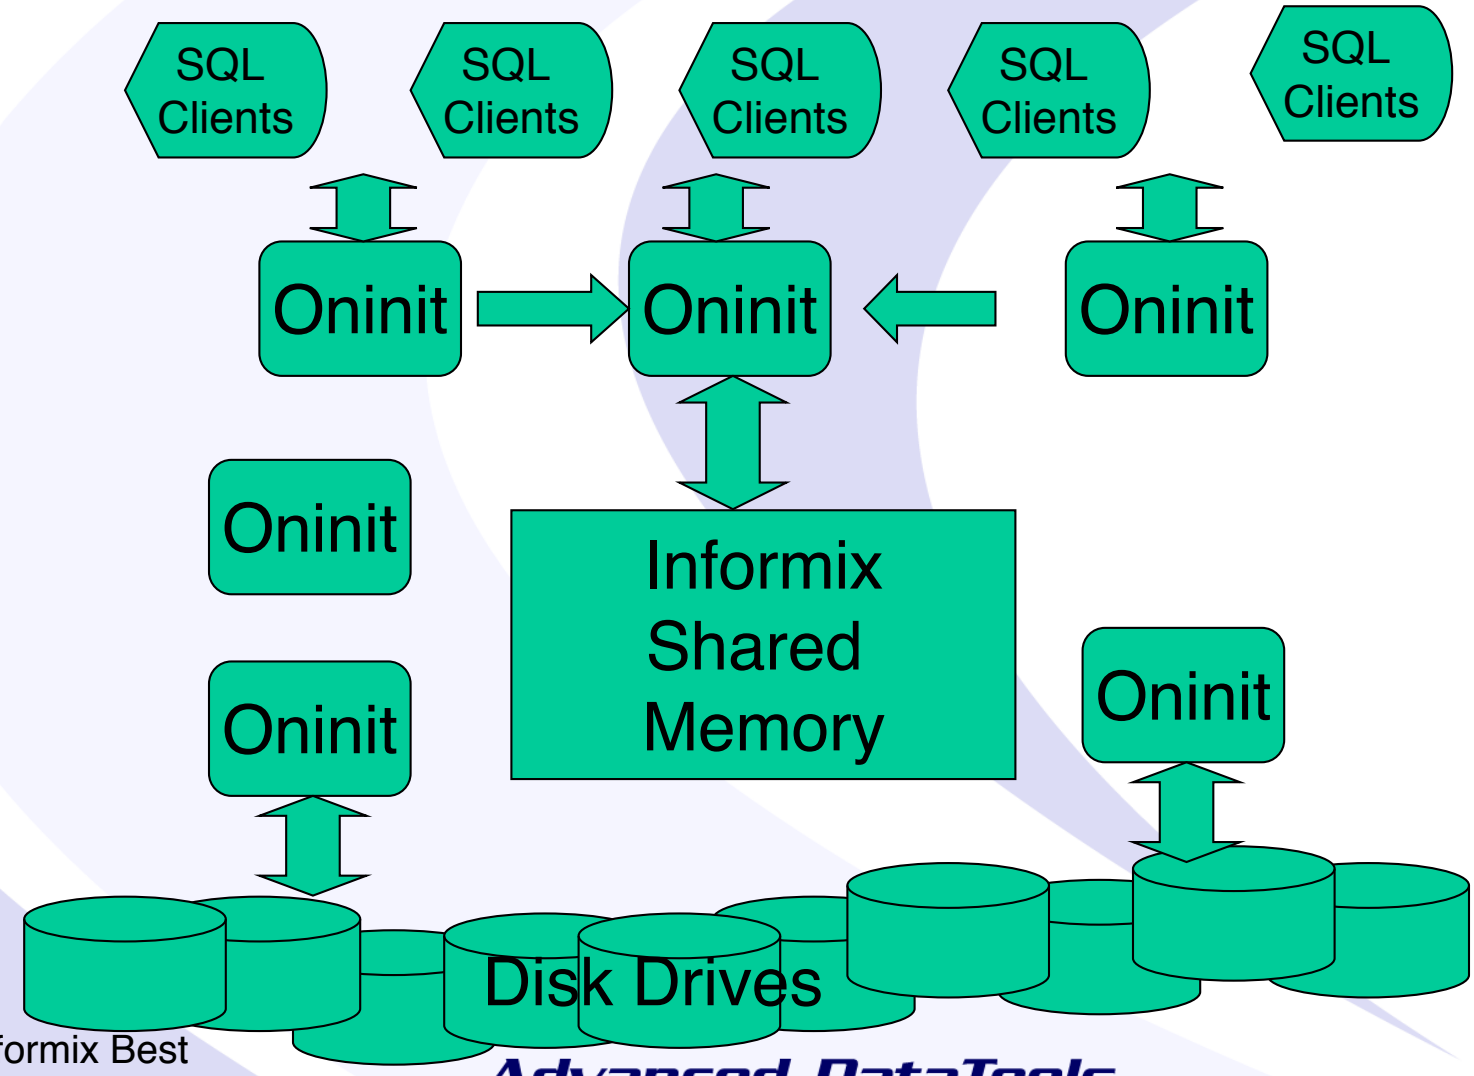

Oninit Process

```
informix@train6:~ train6 > ps -ef | grep oninit
informix 22472     1  9 14:03 ?        00:00:03 oninit -v
root     22473 22472  0 14:03 ?        00:00:00 oninit -v
root     22474 22473  0 14:03 ?        00:00:00 oninit -v
root     22475 22473  0 14:03 ?        00:00:00 oninit -v
root     22476 22473  0 14:03 ?        00:00:00 oninit -v
root     22477 22473  0 14:03 ?        00:00:00 oninit -v
root     22478 22473  0 14:03 ?        00:00:00 oninit -v
root     22479 22473  0 14:03 ?        00:00:00 oninit -v
```

Oninit Process
onstat –g sch

```
informix@train6:~ train6 > onstat -g sch

IBM Informix Dynamic Server Version 12.10.FC6 -- On-Line -- Up 00:02:12 -- 766404 Kbytes
VP Scheduler Statistics:
 vp    pid       class       semops    busy waits  spins/wait  bsy lspins
 1     22472     cpu         141       0           0           0
 2     22473     adm         0         0           0           0
 3     22474     lio         4115      0           0           0
 4     22475     pio         30        0           0           0
 5     22476     aio         7453      0           0           0
 6     22477     msc         5         0           0           0
 7     22478     fifo        2         0           0           0
 8     22479     soc         2         0           0           0
 9     22480     aio         2890      0           0           0
 10    22481     aio         187       0           0           0
 11    22482     aio         113       0           0           0
 12    22483     aio         55        0           0           0
 13    22484     aio         58        0           0           0
 14    22485     aio         41        0           0           0
 15    22486     aio         32        0           0           0
 16    22487     aio         29        0           0           0
 17    22489     aio         22        0           0           0
```

Oninit Process Classes

- CPU - Executes all user and session threads and some system threads
- PIO - Handles physical log file when cooked disk space is used
- LIO - Handles logical log file when cooked disk space is used
- AIO - Handles disk I/O
- SHM - Performs shared memory communications
- TLI - Performs TLI network communications
- SOC - Performs socket network communications
- FIFO - Performs FIFO operations
- OPT - Handles optical disk I/O
- ADM - Executes administrative threads
- ADT - Executes auditing threads
- MSC - Handles request for system calls

Oninit Process Automatically Started

- Started Automatically
  - PIO - Handles physical log file when cooked disk space is used
  - LIO - Handles logical log file when cooked disk space is used
  - FIFO - Performs FIFO operations
  - ADM - Executes administrative threads
  - MSC - Handles request for system calls
- Started when Auditing is on
  - ADT - Executes auditing threads
- Started when UDRs are called
  - Java
  - User Defined Functions

Oninit Process Controlled by VPCLASS

- ONCONFIG VPCLASS Setting
  - CPU - Executes all user and session threads and some system threads
  - AIO - Handles disk I/O

VPCLASS Options

The VPCLASS parameter allows you to:

- Designate a class of virtual processors (VPs)
- Create a user-defined VP, and specify the following information for it:
  - The number of virtual processors that the database server should start initially - optional
  - The maximum number of virtual processors allowed for this class - optional
  - The assignment of virtual processors to CPUs if processor affinity is available - optional
  - The disabling of priority aging by the operating system if the operating system implements priority aging - optional

Syntax:

- VPCLASS classname, options

CPU Oninit Configuration

- VPCLASS CPU – Configure the number of Oninit CPU VPs to start for Informix
  - VPCLASS cpu,num=<number> [,max=<max number cpu>] [,aff=<single CPU number> | <start cpu>-<end cpu> | (<start cpu>-<end cpu>/<skip amount>)]] [,noage]
- Examples for 8 Core machine:
  - VPCLASS cpu,num=4,noage
  - VPCLASS cpu,num=8,noage
  - VPCLASS cpu,num=8,aff=0,noage
  - VPCLASS cpu,num=8,aff=1-4,noage

CPU Affinity

- Example:
  - VPCLASS cpu,num=4,aff=0-3,noage

Message in the Online Log:

15:33:12 Affinitied VP 8 to phys proc 1

15:33:12 Affinitied VP 9 to phys proc 2

15:33:12 Affinitied VP 10 to phys proc 3

15:33:12 Affinitied VP 1 to phys proc 0

Additional CPU Best Practices

- Set MULTIPROCESSOR to 1 (Almost all machines today are multiprocessor)
- Set SINGLE_CPU_VP to 0 (Allows you to run more Oninits of CPU class as needed)
- Set NOAGE if your OS supports it

Additional CPU Best Practices

- Set VP_MEMORY_CACHE_KB <size in KB for private cache for each CPU VP>
- Format is: <size>[,DYNAMICISTATIC] Acceptable values for <size> are: 0 (disable) or 800 through 40% of the value of SHMTOTAL
- Example:
  - VP_MEMORY_CACHE_KB 4096

AIO Oninit Best Practices

- Default is
  - AUTO_AIOVPS 1 – enable automatically adding AIO VPs as needed
  - This can lead to ***too many*** AIO VPs writing to the same disk system
- Recommended
  - AUTO_AIOVPS 0
  - VPCLASS aio,num=<number of oninits you need to write to disk>

AIO Oninit Best Practices

- How many AIO Class Oninits do you need? Test, Test, Test…
  - With KAIO on – only need 2 AIO oninits
  - With KAIO off (default), it depends on how many processes can write to a disk at the same time.
  - Never need more than twice the number of active chunks.
  - Most disks can handle up to 8 AIO oninit processes.

AIO Oninit Best Practice Examples

- 1 Disk and 24 Chunks
  - VPCLASS aio,num=8
- 6 Disks and 24 Chunks (12 active)
  - VPCLASS aio,num=24

Oninit Processes Controlled by NETTYPE

- ONCONFIG NETTYPE Setting
  - SHM - Performs shared memory communications
  - TLI - Performs TLI network communications
  - SOC - Performs socket network communications

NETTYPE Configuration

- NETTYPE <protocol>,<number of oninit process>, <number of connections per oninit>, <Type of Oninit – CPU or NET>
- Examples:
  - NETTYPE ipcshm,1,50,CPU
  - NETTYPE soctcp,4,250,NET

NETTYPE Configuration Best Practices

- Configure Shared Memory Connection to run on NETTYPE type CPU and Network Connections to run on NETTYPE NET.
- Configure 200-300 Connections per Oninit process.
- Example:
  - NETTYPE ipcshm,1,50,CPU – Shared Memory with 50 connections
  - NETTYPE soctcp,4,250,NET – Network with 1000 connections

Memory Recommendations for Informix and Best Practices

Advanced DataTools

Informix Memory Best Practices

- How much memory is available on the machine?
- How much is used by the Operating System and other applications?
- How much will be assigned to Informix?
- ***DO NOT allow the machine to Swap memory to disk as this will SLOW everything down***

Informix Memory Classes

- R – Resident Memory Segment
- B – Buffer Pool Segment for data
- V – Virtual Memory Segment for Working Storage
- M – Message Segment for communications between clients

Informix Shared Memory

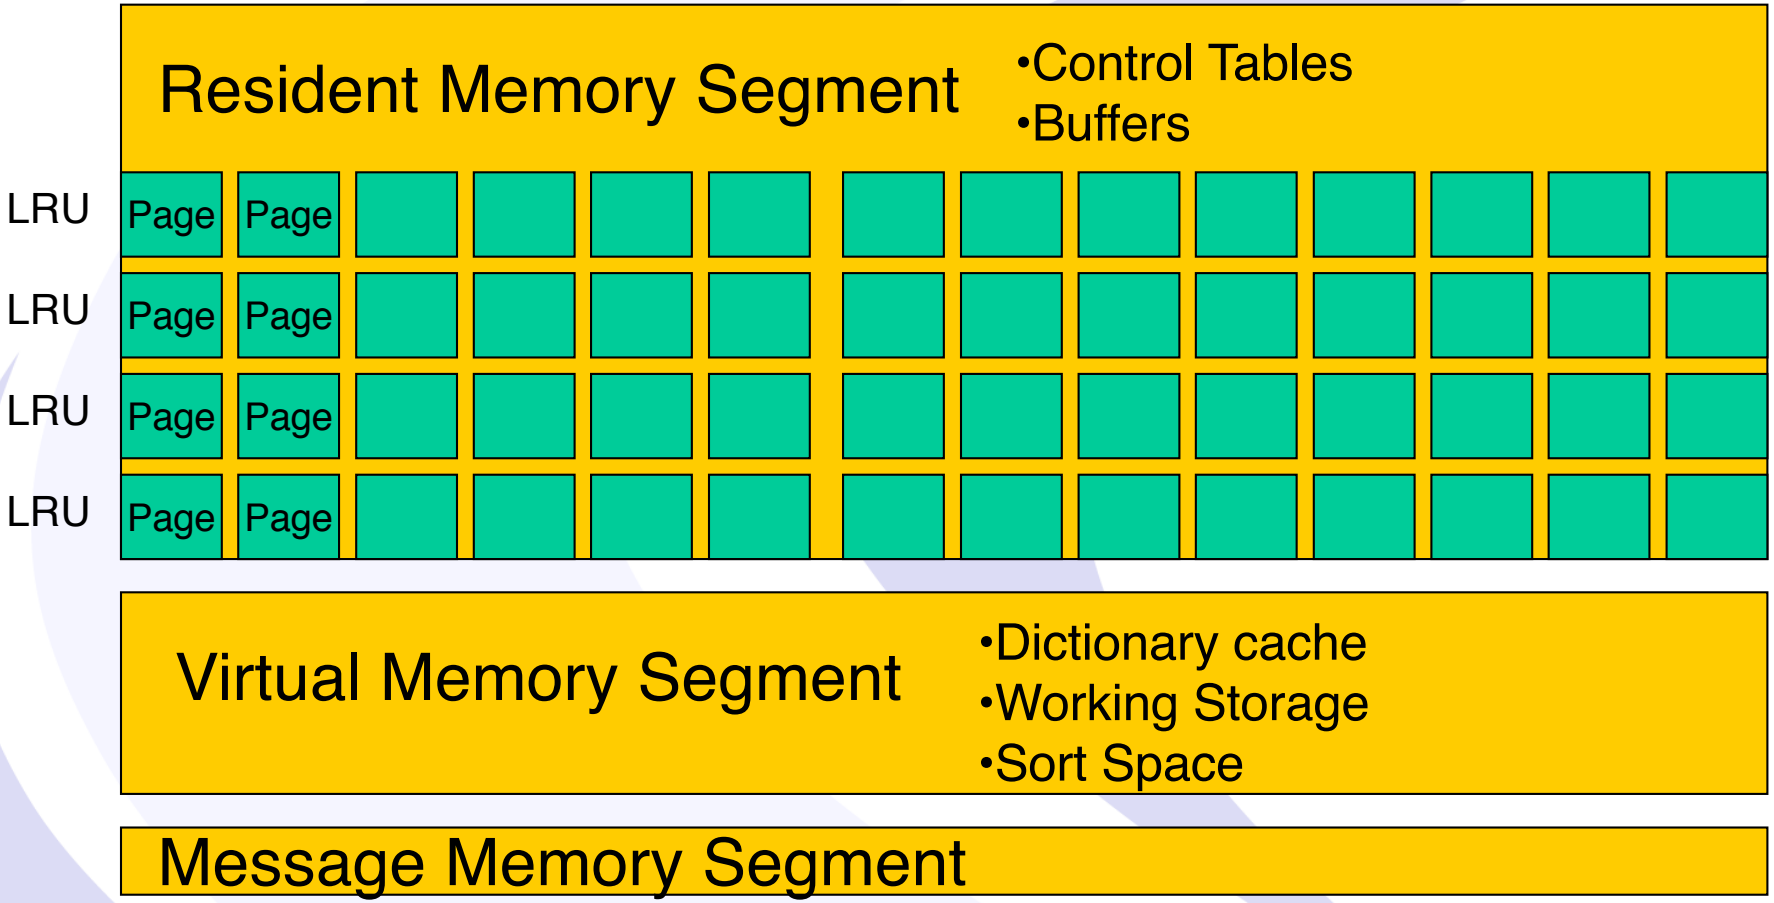

Informix Shared Memory onstat –g seg

```
informix@train6:~ train6 > onstat -g seg

IBM Informix Dynamic Server Version 12.10.FC6 -- On-Line -- Up 00:30:00 -- 766404 Kbytes

Segment Summary:
id          key         addr              size              ovhd     class blkused  blkfree
32769       525c4801    44000000          4911104           495784   R     1199     0
65538       525c4802    444af000          33439744          393384   V     8030     134
98307       525c4803    46493000          562749440         1        B     137390   0
131076      525c4804    67d41000          166359040         1        B     40615    0
163845      525c4805    71be8000          561152            7848     M     136      1
196614      525c4806    71c71000          8388608           99720    V     1529     519
294919      525c4807    72471000          8388608           99720    V     25       2023
Total:      -           -                 784797696         -        -     188924   2677

   (* segment locked in memory)
No reserve memory is allocated
```

ONCONFIG Memory Configuration Settings

- BUFFERPOOL
- SHMVIRTSIZE
- SHMADD
- LOCKS

BUFFERPOOL Best Practices

- More Buffers - the better and ***faster*** your database will perform.
- Goal is to put all the active data into Memory Buffers.
- Goal is to prevent high Memory Buffers Turnover (Art Kagel's rule – less than 8 times and hour).
- Always leave the default BUFFERPOOL line in the ONCONFIG.

BUFFERPOOL Settings

- The BUFFERPOOL configuration parameter consists of two lines in the onconfig.std file, as shown in this example for a UNIX platform:

```
BUFFERPOOL default,lrus=8,buffers=5000,lru_min_dirty=50,lru_max_dirty=60
BUFFERPOOL size=2K,buffers=5000,lrus=8,lru_min_dirty=50,lru_max_dirty=60
```

- The top line specifies the default values that are used if you create a dbspace with a page size that does not already have a corresponding buffer pool created at startup.

- The next line below the default line specifies the database server's default values for a buffer pool, which are based on the database server's default page size.

- When you add a dbspace with a different page size with the onspaces utility, or when you add a new buffer pool with the onparams utility, a new line is appended to the BUFFERPOOL configuration parameter in the ONCONFIG file. The page size for each buffer pool must be a multiple of the system's default page size.

BUFFERPOOL Examples

- 3 GB Memory for Buffers – Linux OLTP System
  - BUFFERPOOL size=2k,buffers=1500000,lrus=32,lru_min_dirty=10,lru_max_dirty=20
- 12 GB Memory for Buffers – AIX OLTP System
  - BUFFERPOOL size=4k,buffers=3000000,lrus=128,lru_min_dirty=1,lru_max_dirty=2
- 48 GB Memory for Buffers – Solaris Data Warehouse
  - BUFFERPOOL size=2K,buffers=24000000,lrus=128,lru_min_dirty=60,lru_max_dirty=70
- 15 GB Memory for 4K Buffers and 12.8 GB for 16K Buffers
  - BUFFERPOOL size=4K,buffers=60000000,lrus=256,lru_min_dirty=0.1,lru_max_dirty=0.2
  - BUFFERPOOL size=16K,buffers=800000,lrus=256,lru_min_dirty=20,lru_max_dirty=30

Monitoring BUFFERPOOL Turnover

```
--------------------------------------------------------------------------------
-- Module: @(#)buff_btr_ratio.sql     2.0     Date: 2013/04/10
-- Author: Lester Knutsen  Email: lester@advancedatatools.com
--         Advanced DataTools Corporation
-- Discription: Display Buffer Turnovers per hour
--              Based on Art Kagels performance tuning tip on monitoring
--              how much buffer churn your server has.
--              Goal is BTR of less then 7 times per hour
--      Tested with Informix 11.70 and Informix 12.10
--------------------------------------------------------------------------------

select
        bufsize,
        pagreads,
        bufwrites,
        nbuffs,
        ((( pagreads + bufwrites ) /nbuffs ) /
                ( select (ROUND ((( sh_curtime - sh_pfclrtime)/60)/60) )
                from sysshmvals ) ) BTR
from sysbufpool;
```

Informix Shared Memory Management

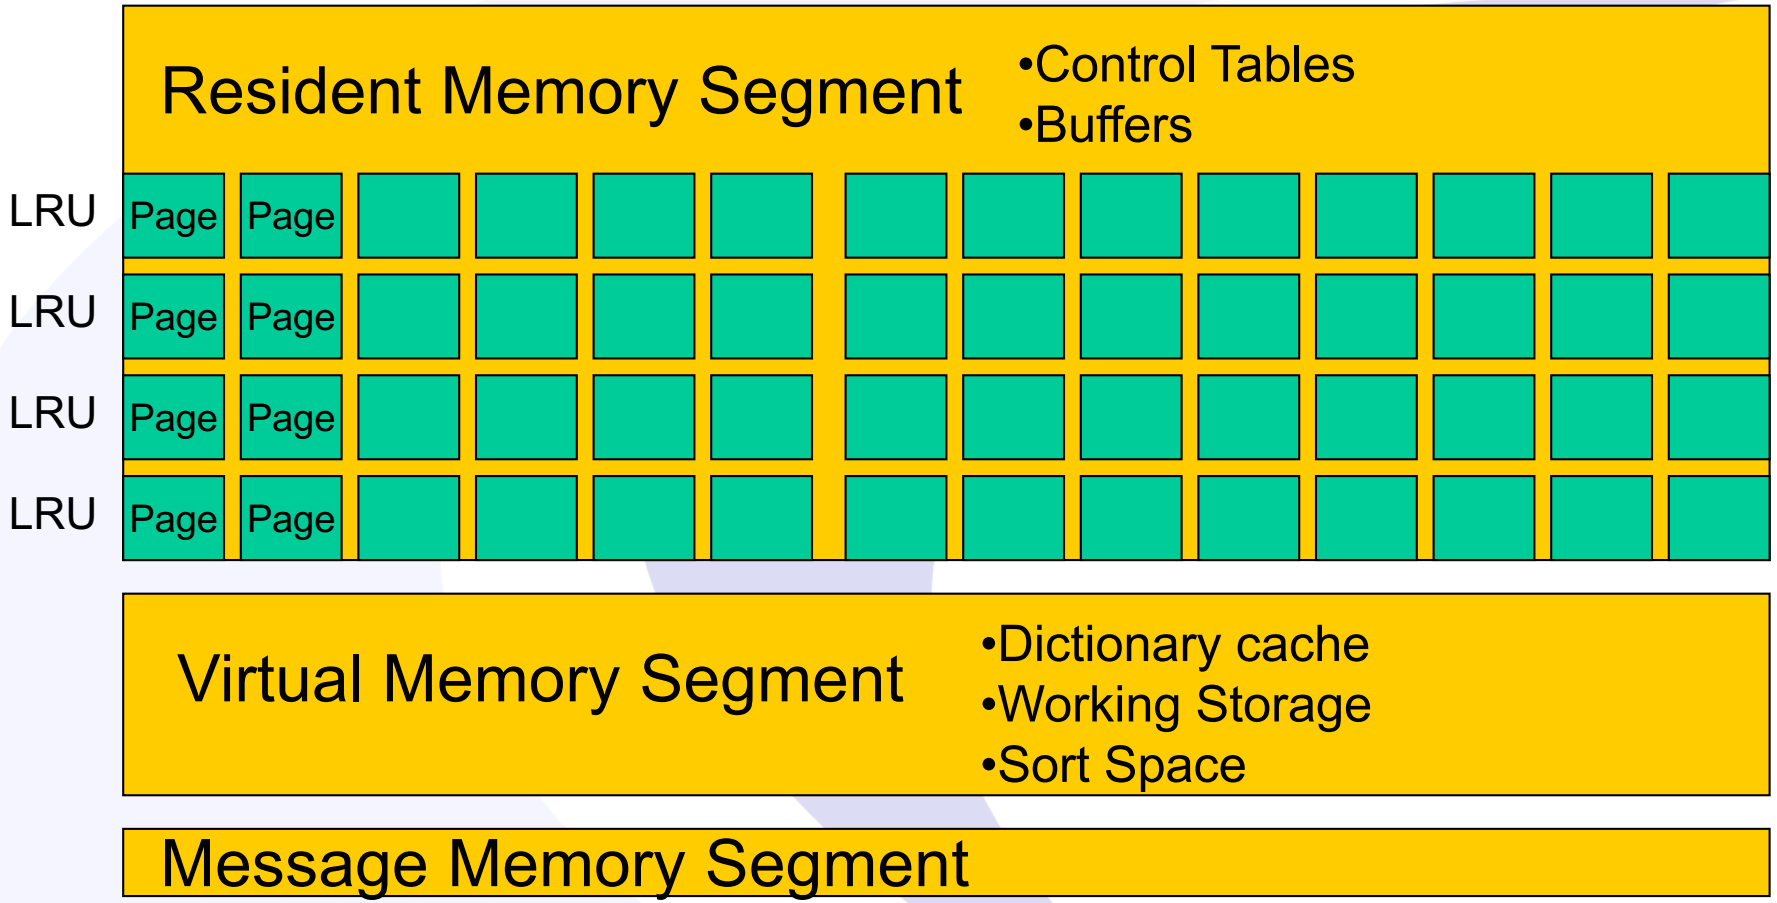

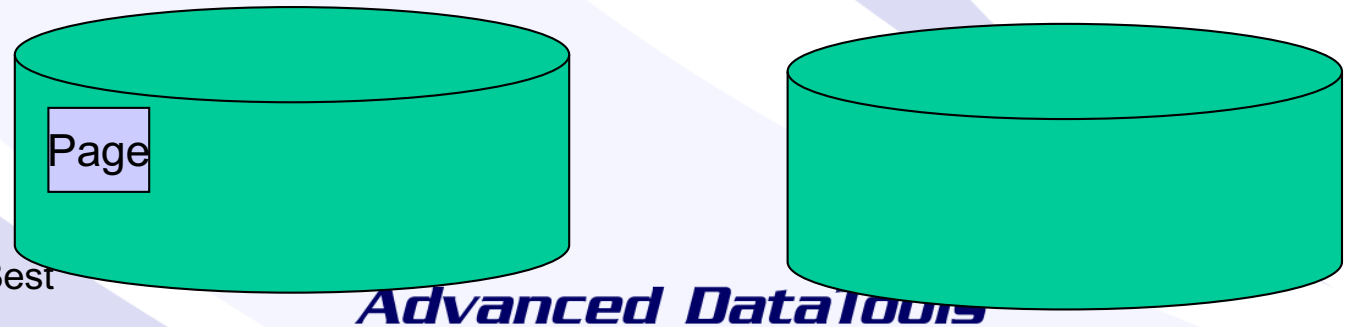

Page Gets Read into Memory by a Select

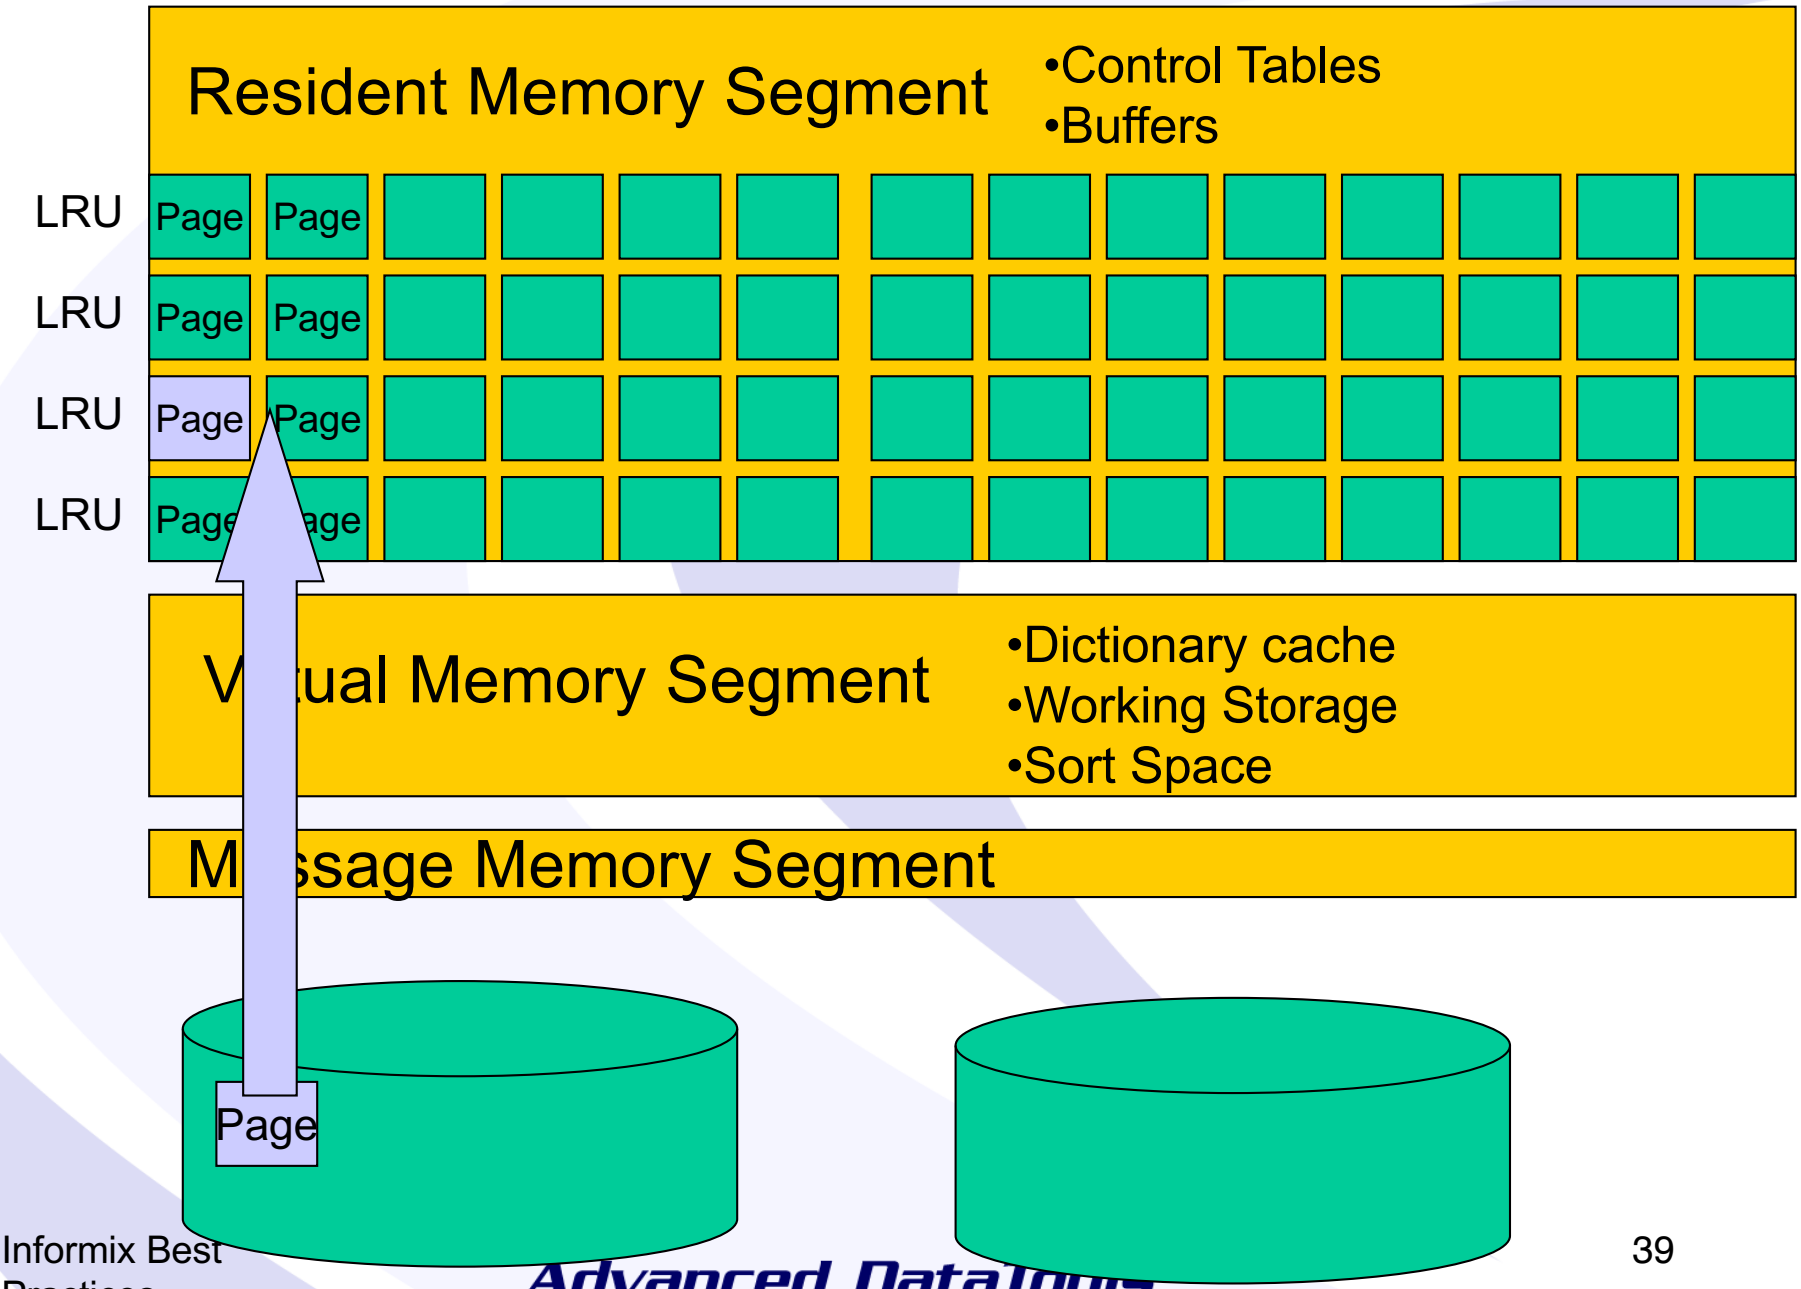

When all Buffers are full, Least Recently Used (LRU) Page is discarded to make room for more data.

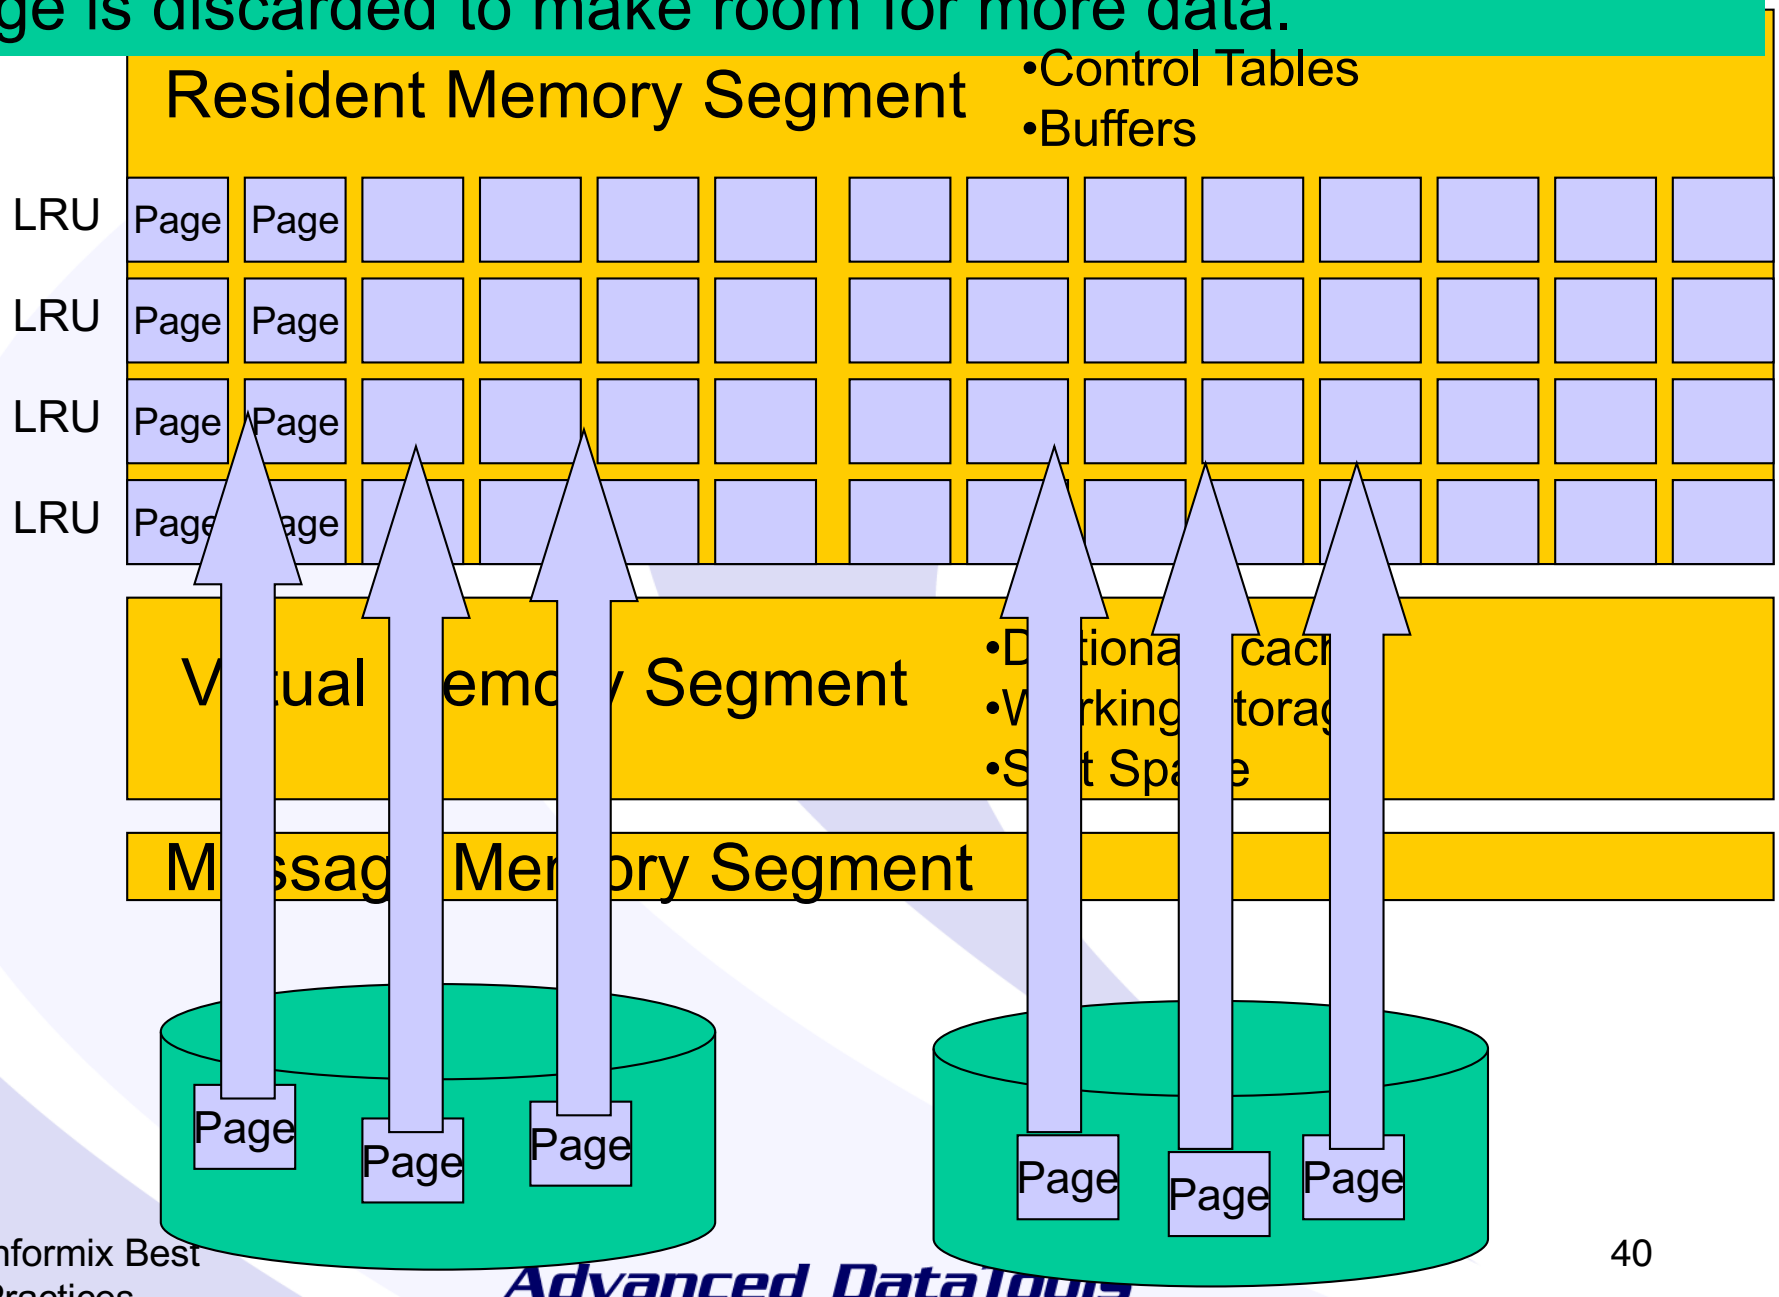

When a user updates a page, it is marked as dirty and must be written out to disk before it is discarded.

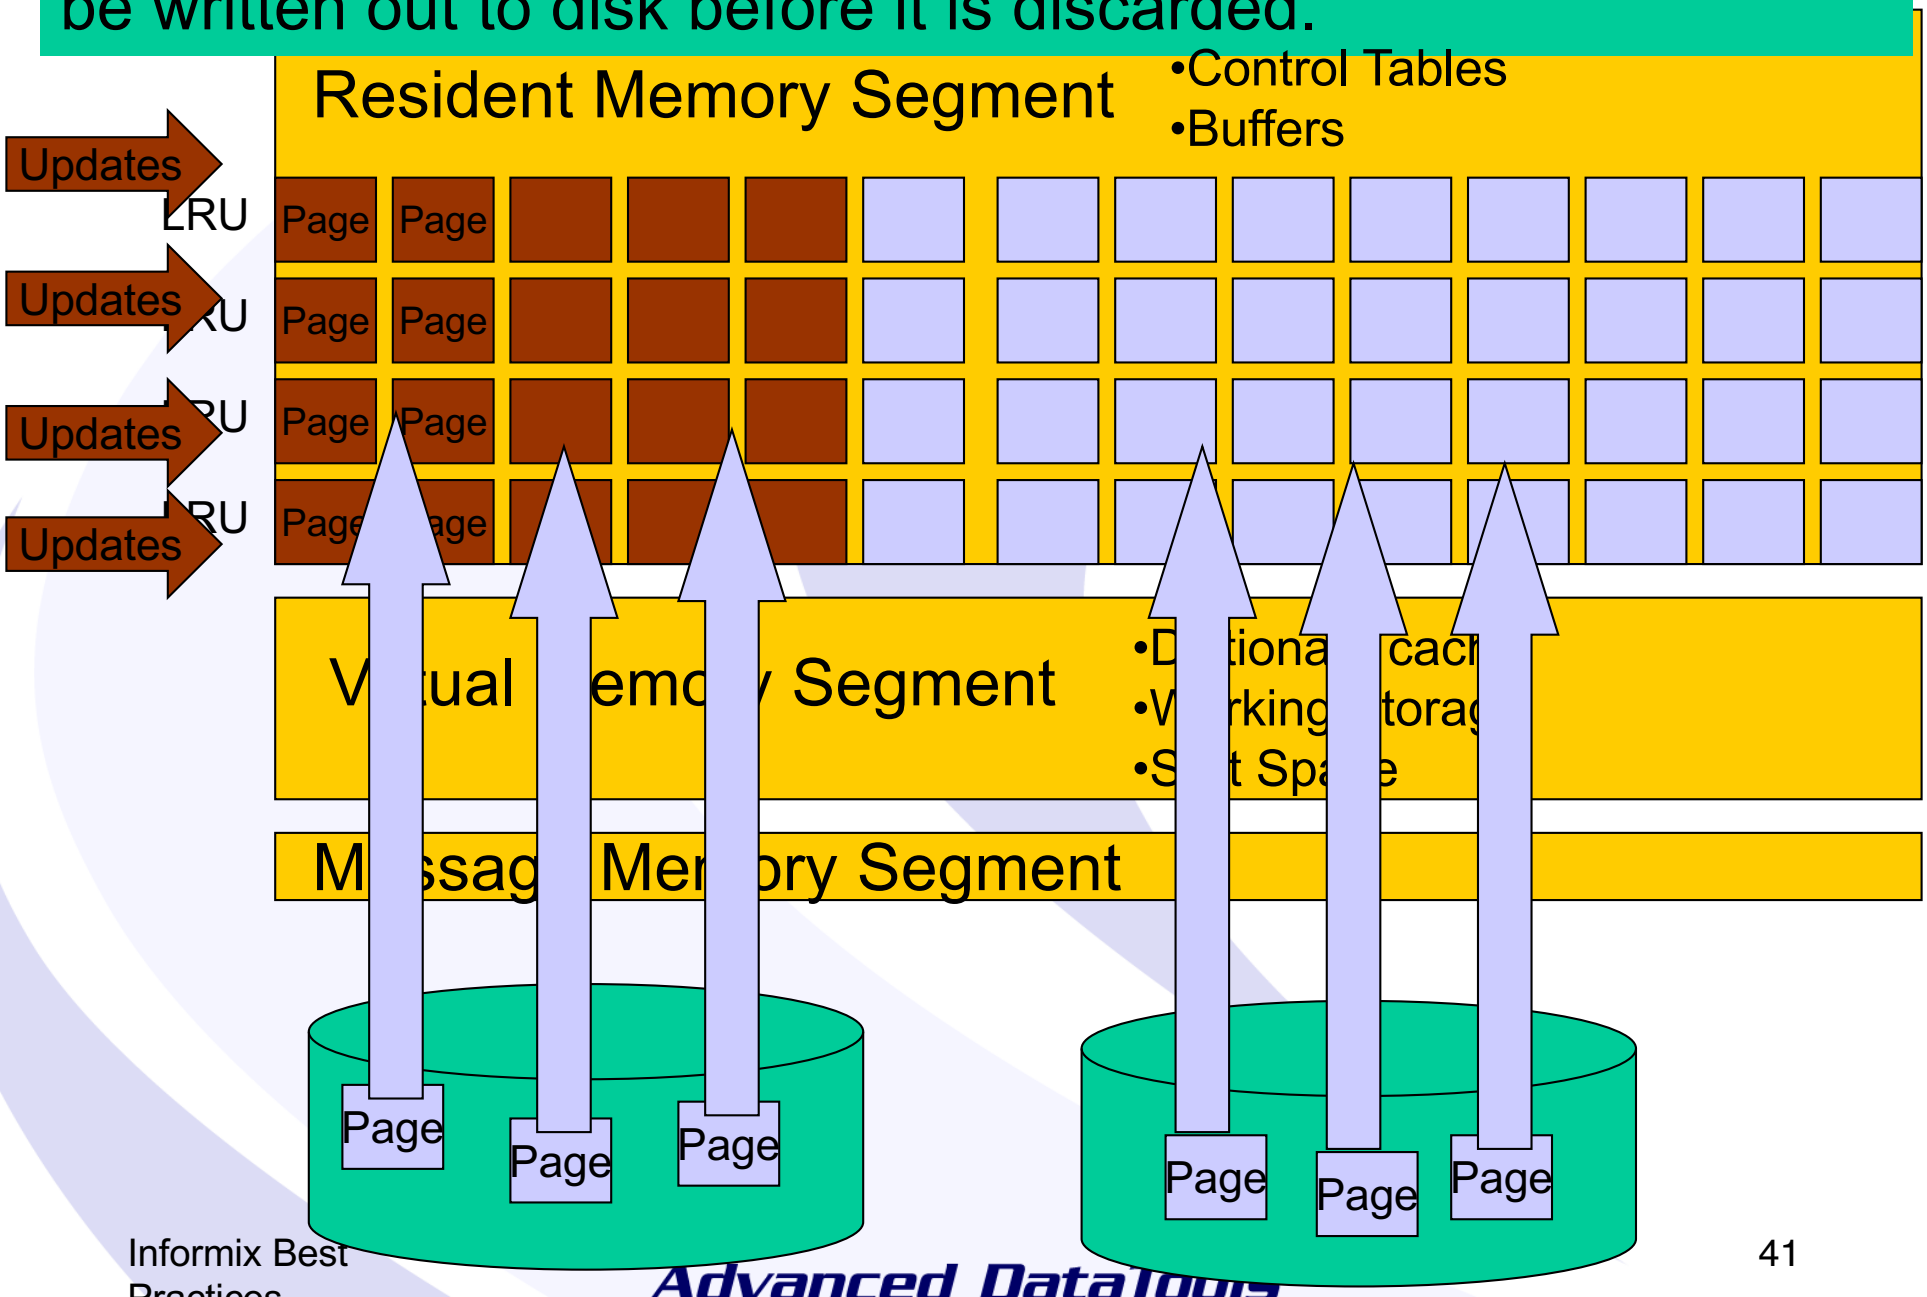

LRU Queues manage writing data to disk in the background when there is idle time based on LRU_MAX_DIRTY and LRU_MIN_DIRTY ONCONFIG values.

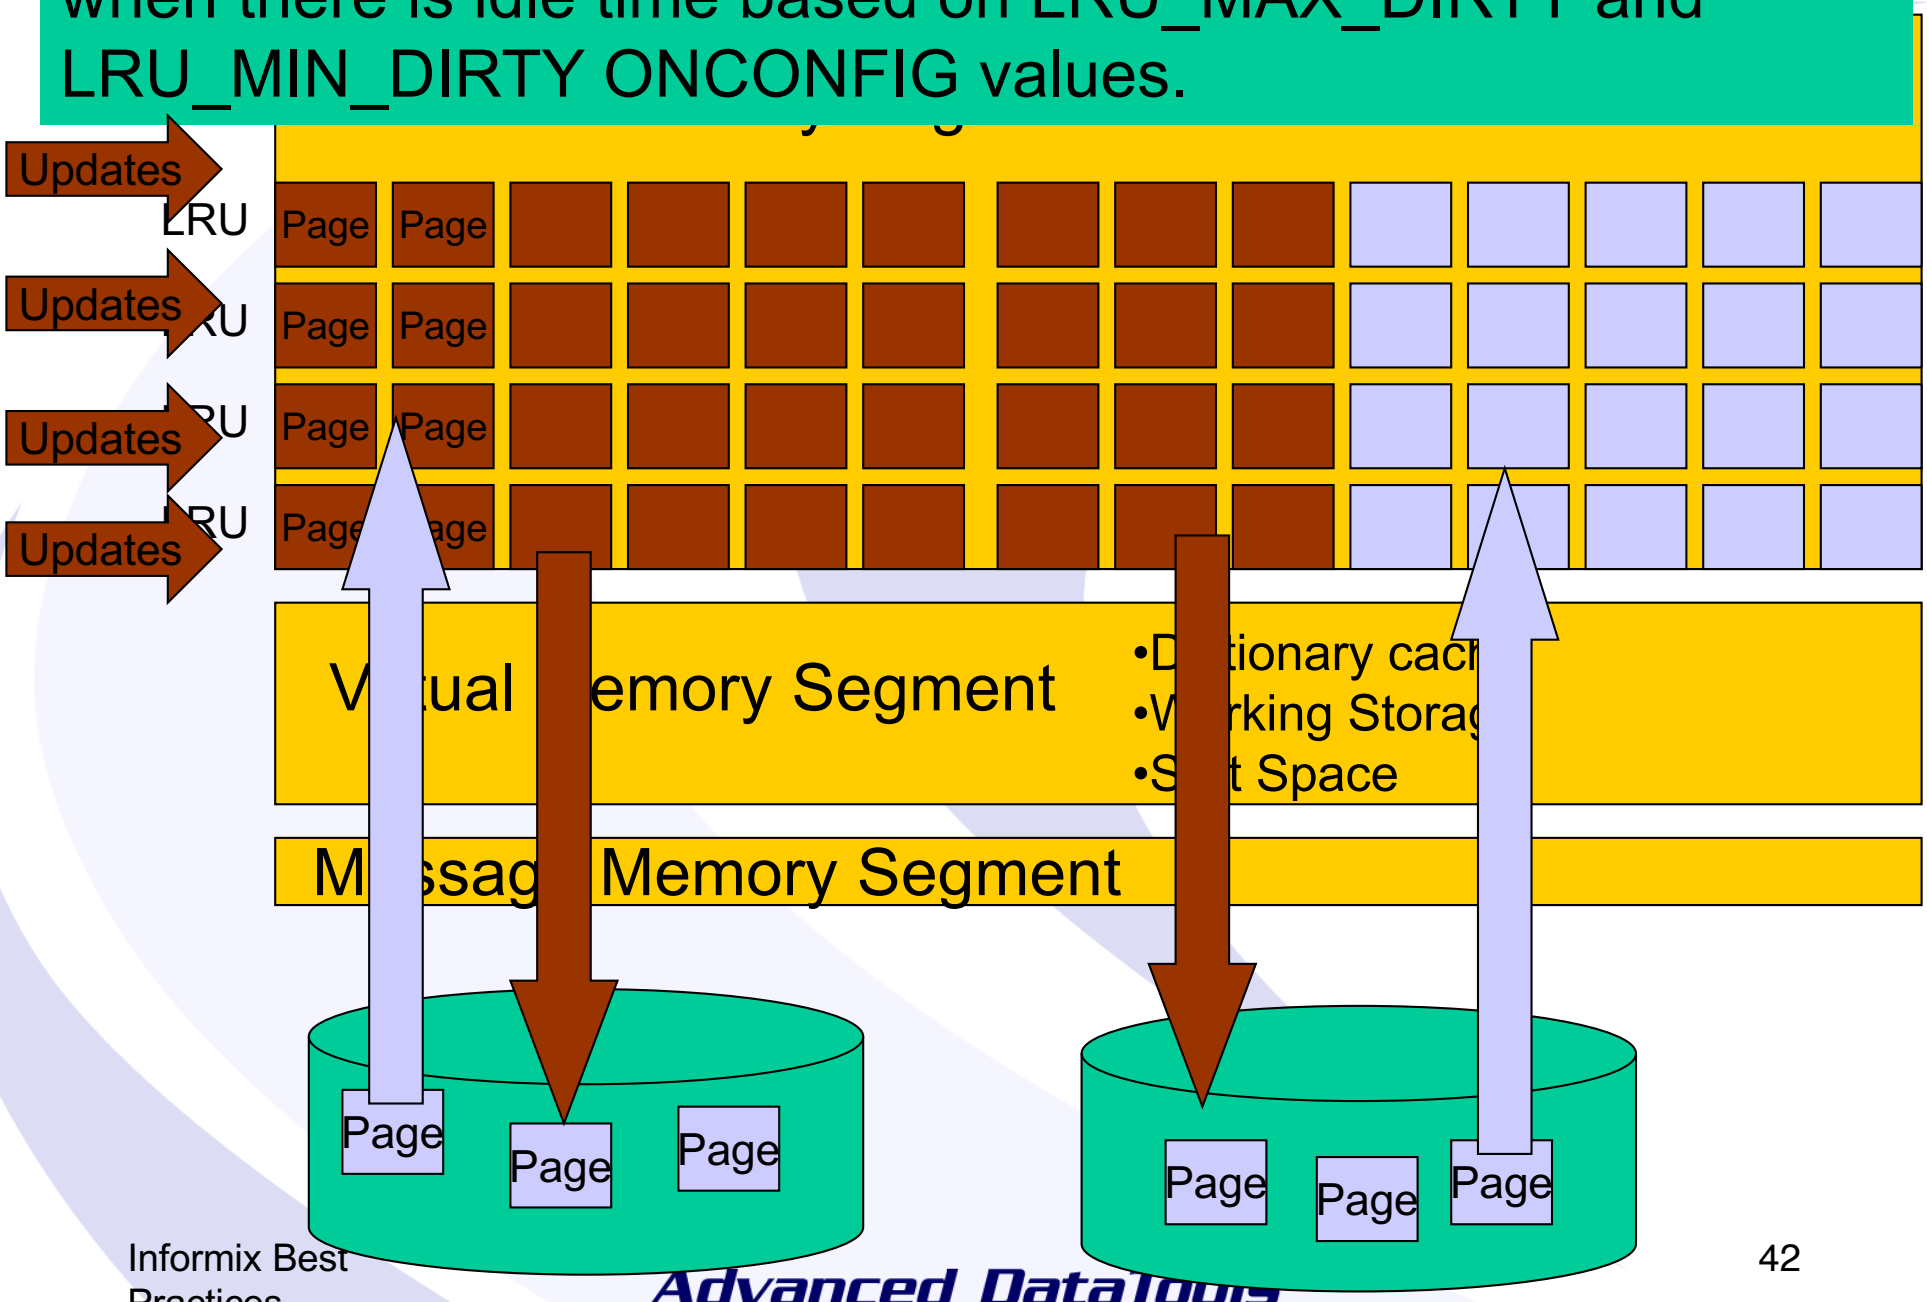

When all buffers are Dirty, the server must STOP all processing and perform a Foreground Write.

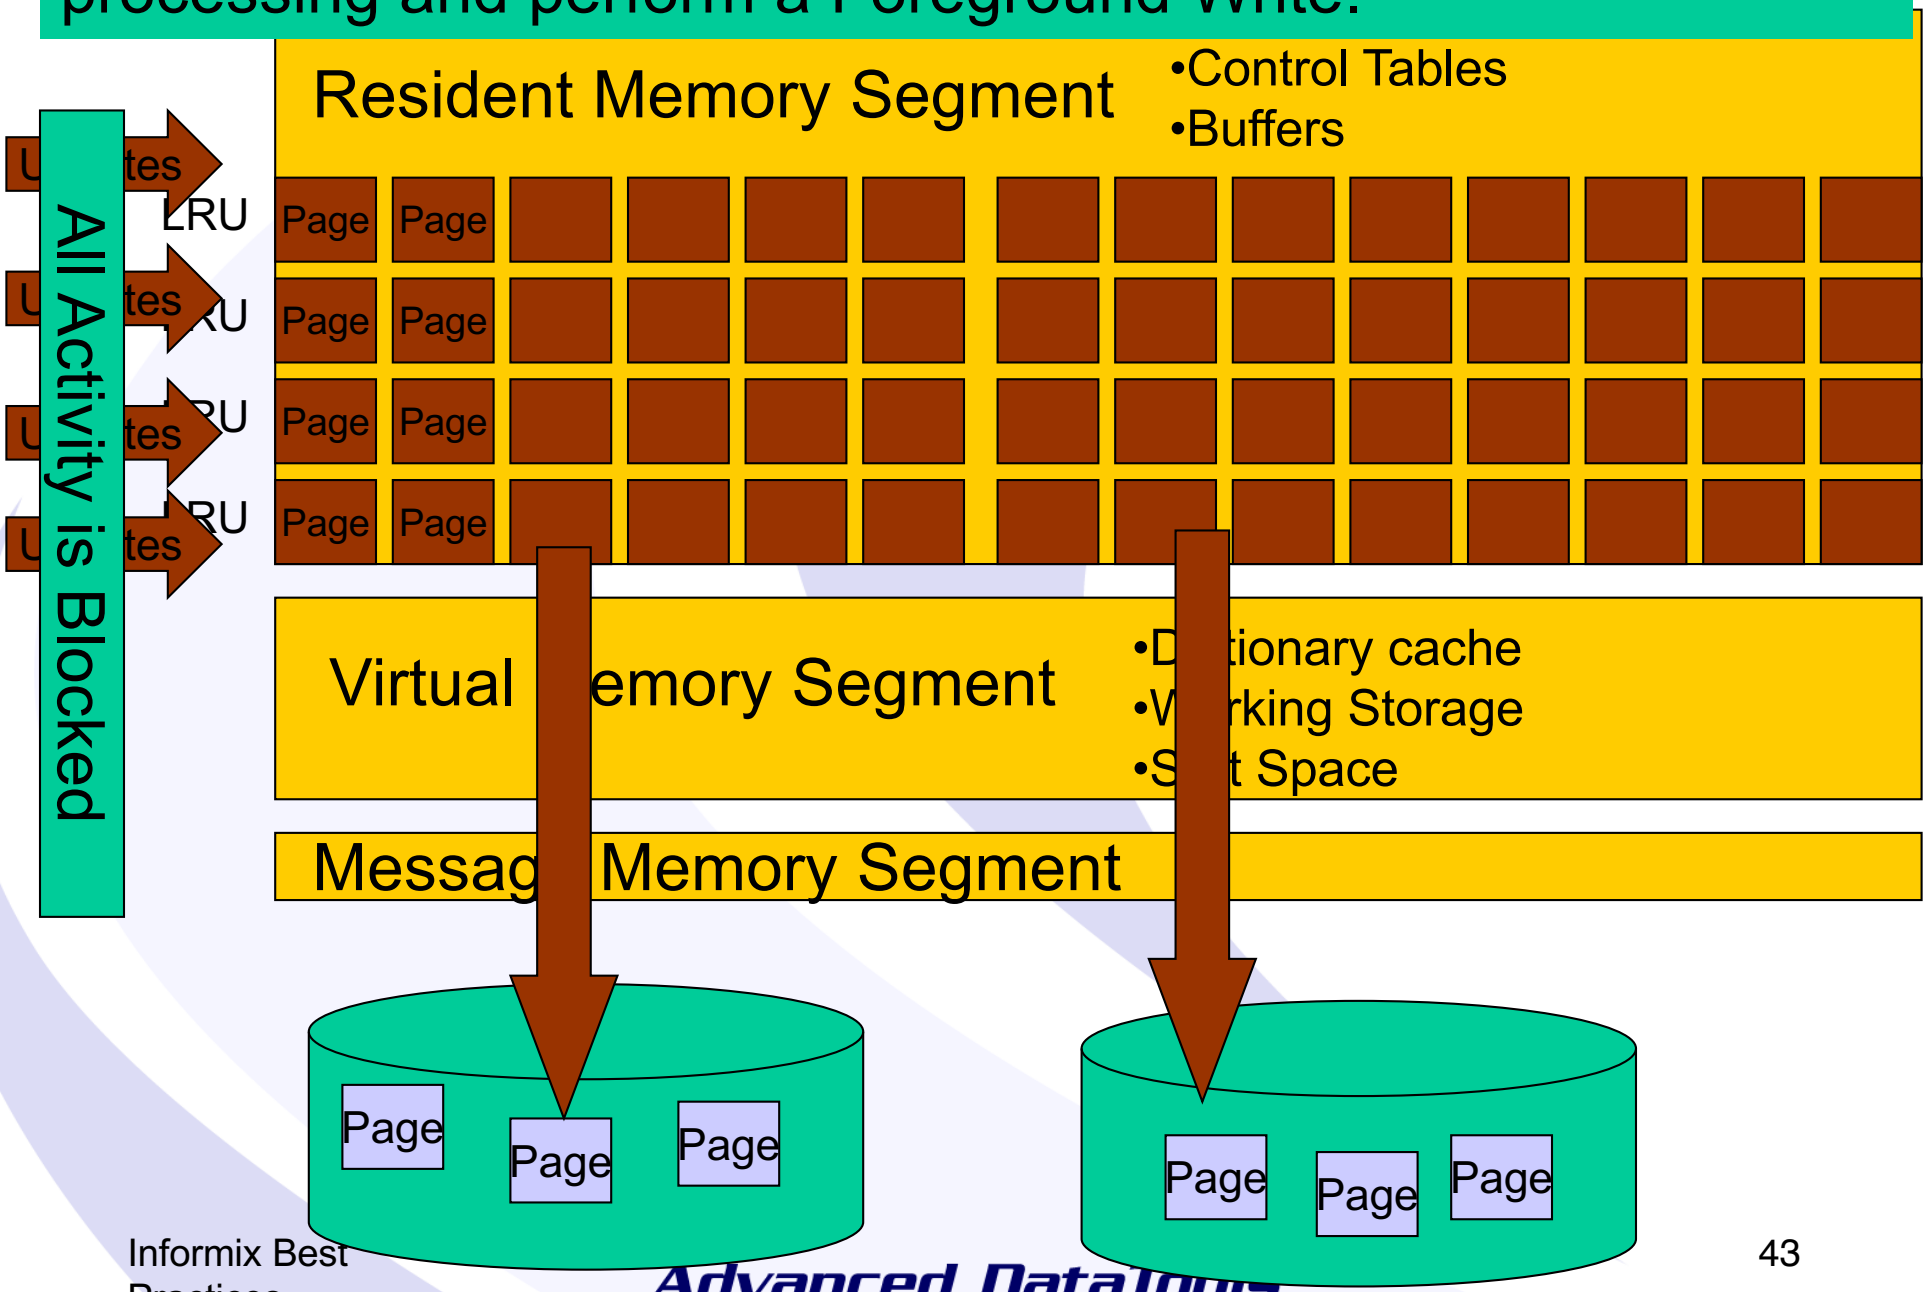

When a buffer is written to disk, it is marked as clean and may be discarded if needed.

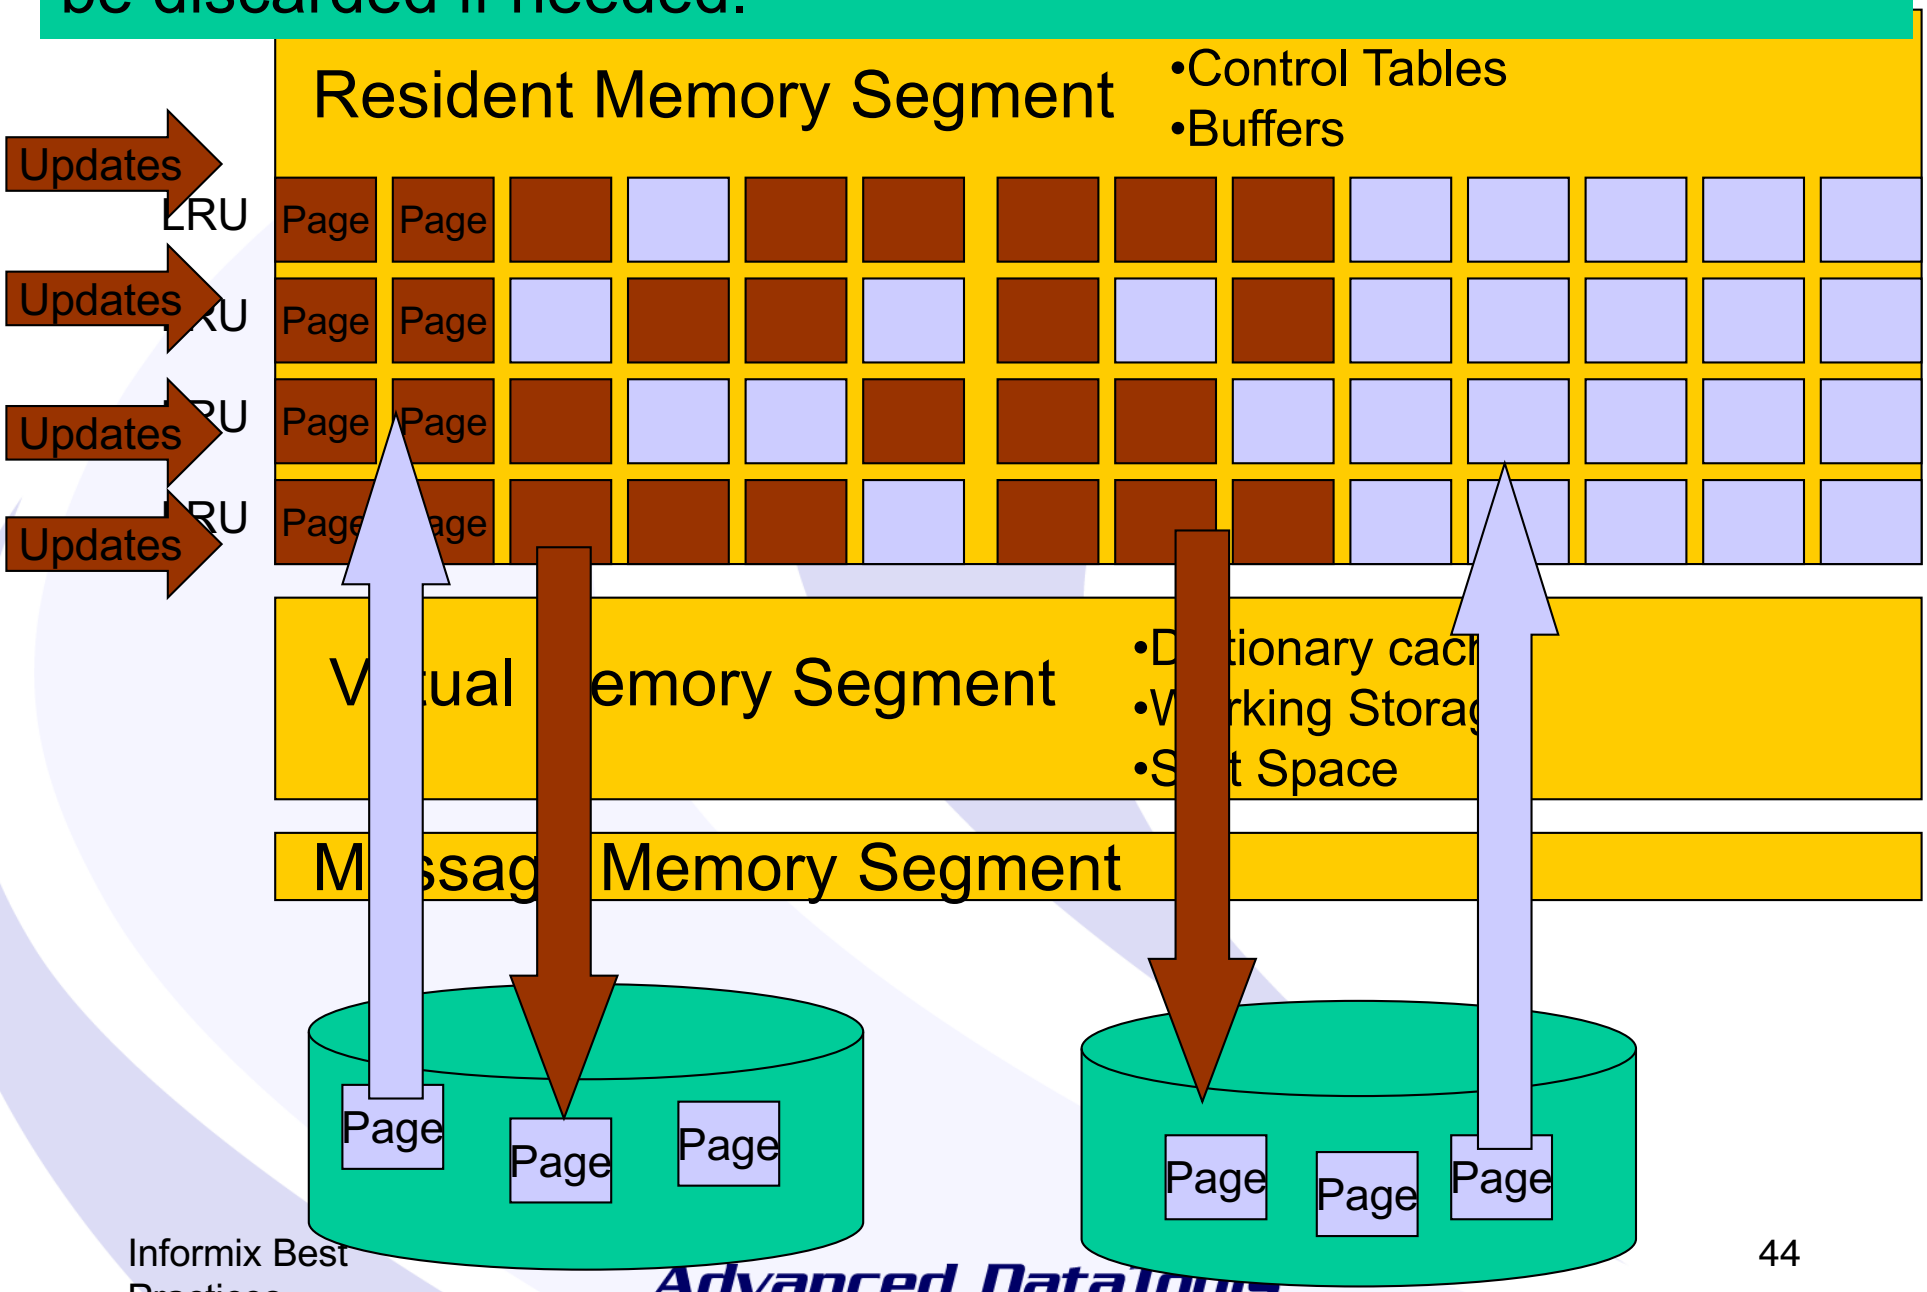

Checkpoint writes all Dirty Buffers to Disk.

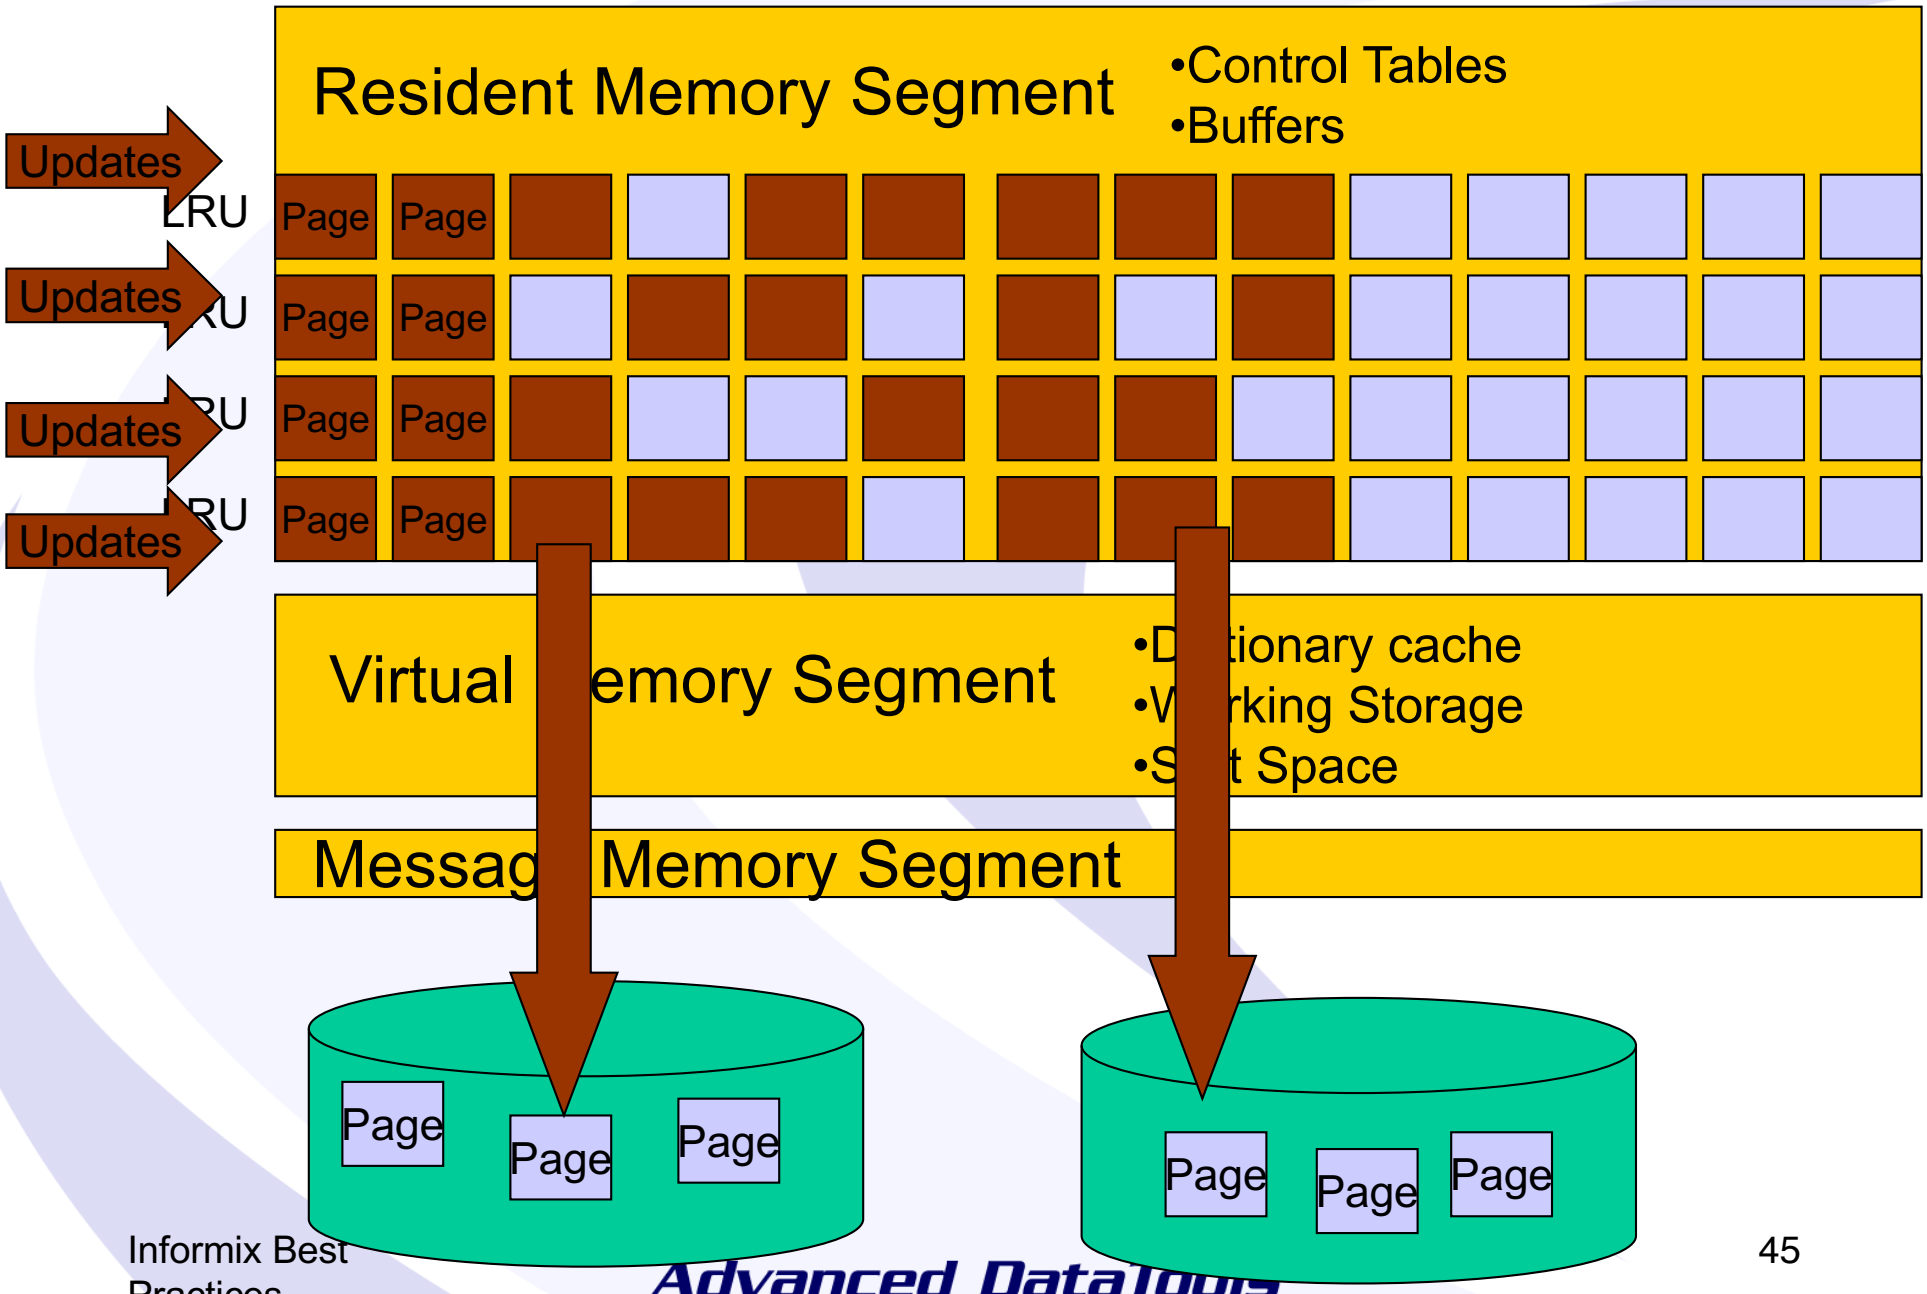

After a Checkpoint, all Buffers are clean and the cycle starts over again.

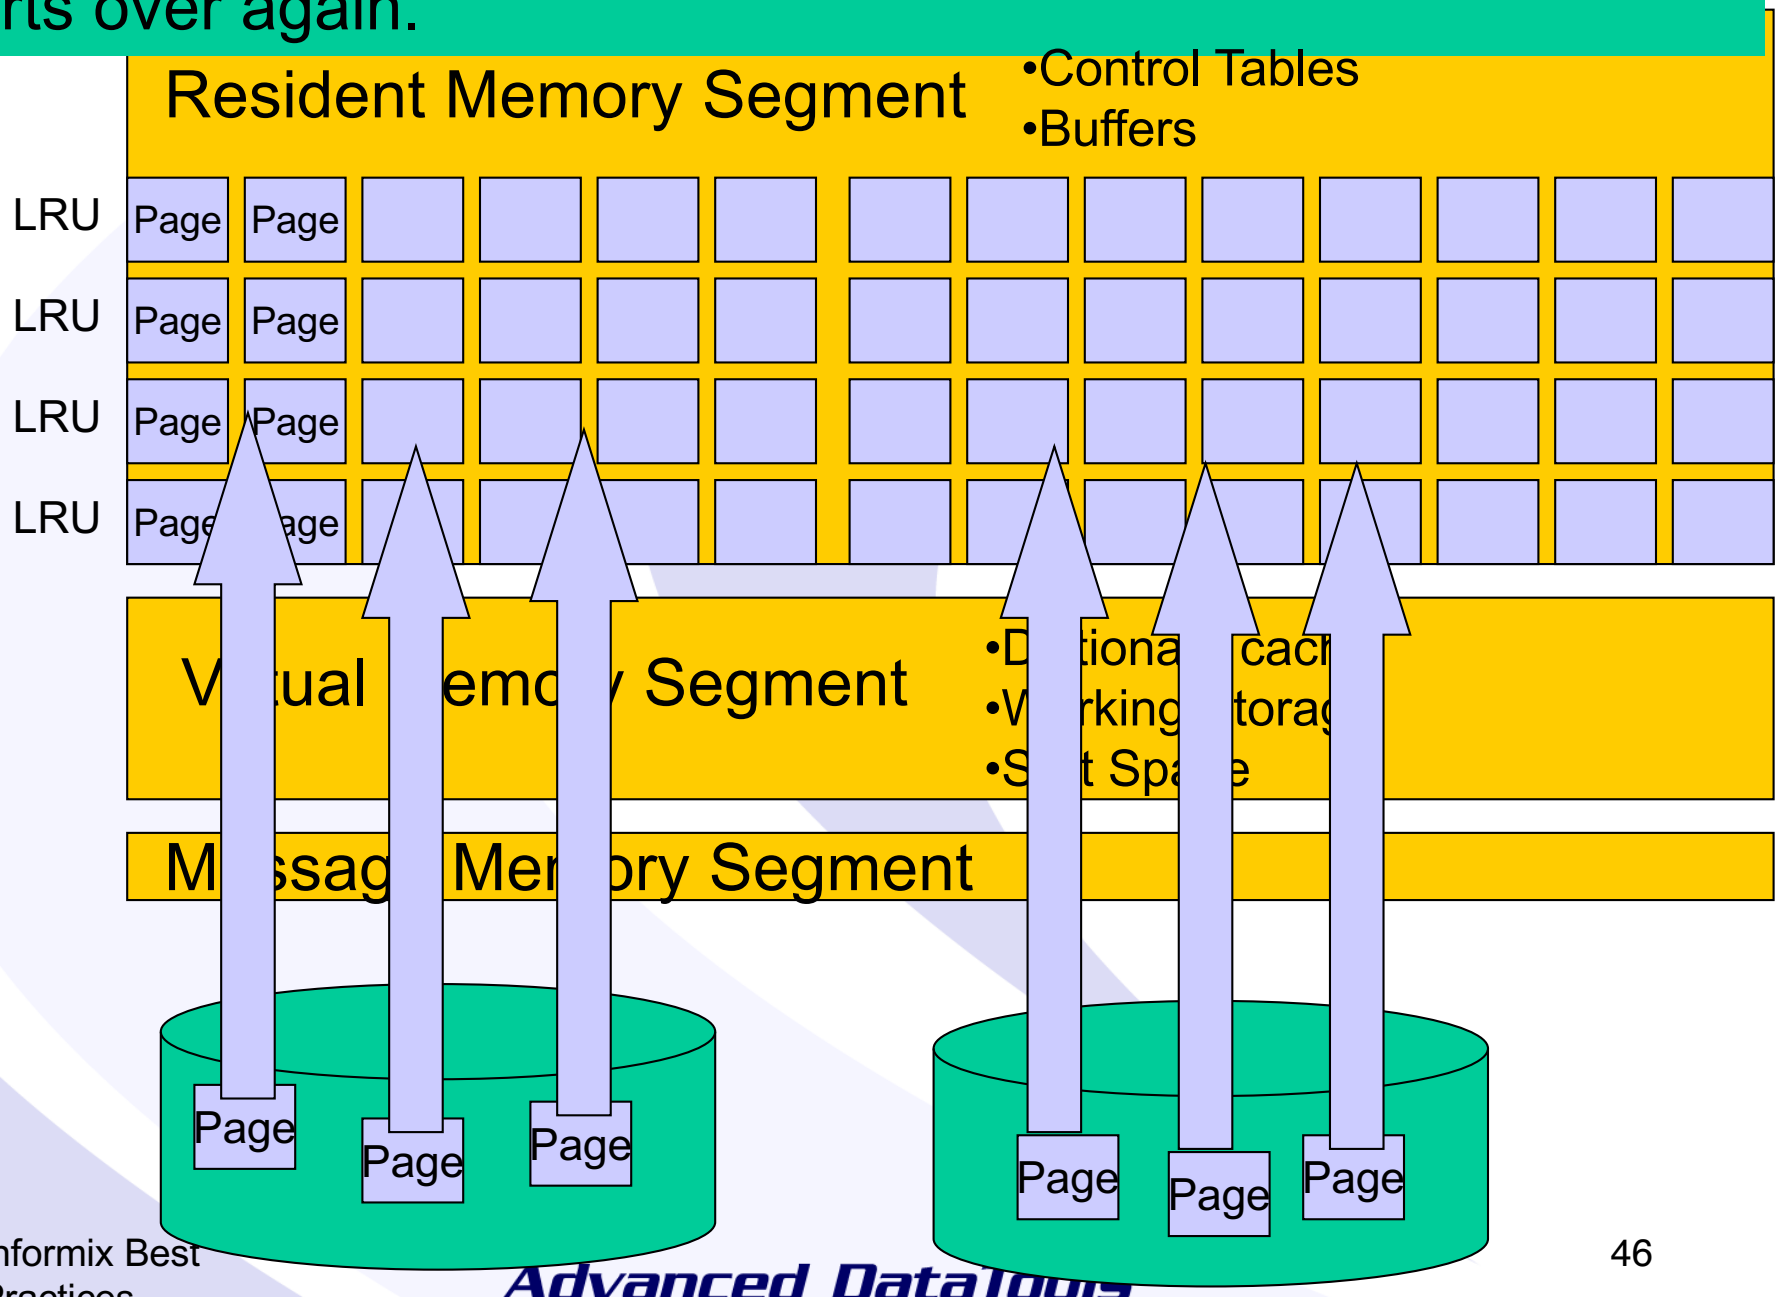

Memory LRU Settings

- AUTO_LRU_TUNING - Enables (1) or disables (0)
- BUFFERPOOL LRU Settings
  - lrus=<Number of LRU QUEUES>,
  - lru_max_dirty=<Percent dirty to START cleaning>
  - lru_min_dirty=<Percent dirty to STOP cleaning>

Memory LRU Best Practices

- Enable AUTO_LRU_TUNING for turnkey or embedded systems.
- Disable AUTO_LRU_TUNING for high performance systems where you do not want CHECKPOINTS to write a huge amount of data to disk and slow everything down.

Memory LRU Settings

- LRU Settings for AUTO_LRU_TUNING Disabled
  - lrus=<Number of LRU QUEUES>,
  - lru_max_dirty=<Percent dirty to START cleaning>
  - lru_min_dirty=<Percent dirty to STOP cleaning>

SHMVIRTSIZE Best Practices

- Controls the size of the Informix Virtual Memory Workspace, which can grow if needed.
- Best practice is to set it large enough so it does not need to grow.
  - Monitor with onstat –g seg

Informix Shared Memory onstat –g seg

```
informix@train6:~ train6 > onstat -g seg

IBM Informix Dynamic Server Version 12.10.FC6 -- On-Line -- Up 00:30:00 -- 766404 Kbytes

Segment Summary:
id          key         addr              size              ovhd     class blkused  blkfree
32769       525c4801    44000000          4911104           495784   R     1199     0
65538       525c4802    444af000          33439744          393384   V     8030     134
98307       525c4803    46493000          562749440         1        B     137390   0
131076      525c4804    67d41000          166359040         1        B     40615    0
163845      525c4805    71be8000          561152            7848     M     136      1
196614      525c4806    71c71000          8388608           99720    V     1529     519
294919      525c4807    72471000          8388608           99720    V     25       2023
Total:      -           -                 784797696         -        -     188924   2677

   (* segment locked in memory)
No reserve memory is allocated
```

Additional Memory Setting

- RESIDENT - Controls whether shared memory is resident. Acceptable values are:
  - 0 off (default)
  - 1 lock the resident segment only
  - n lock the resident segment and the next n-1 virtual segments, where n < 100
  - -1 lock all resident and virtual segments
- SHMADD - The size, in KB, of additional virtual shared memory segments

LOCKS Memory Settings

- LOCKS – The number of LOCKS when Informix Starts. This determines the amount of Memory initially set for LOCKS. Can be dynamically added when needed.
- Dynamically adding LOCKS can cause a performance degradation.

LOCKS Memory Settings

- To monitor, look at the last line of:
  - onstat –k

```
IBM Informix Dynamic Server Version 12.10.FC6 -- On-Line -- Up 02:34:23 -- 165016 Kbytes

Locks
address          wtlist           owner            lklist           type     tblsnum  rowid
44199028         0                44cd4668           0                    S    100002   204
4423f068         0                44cd4f28           0                    S    100002   204
442e50a8         0                44cd57e8           0                    S    100002   204
442e5130         0                44cd57e8           442e50a8         HDR+S    100002   201
4438b0e8         0                44cd60a8           0                HDR+S    100002   204
 5 active, 80000 total, 16384 hash buckets, 2 lock table overflows
```

  - This shows 2 lock table overflows.
  - This system requires 80,000 locks.

LOCK Best Practices

- LOCK Table Overflows will slow performance and should be avoided.
- LOCK Table Overflows are a major contributor to SHMVIRT Memory additions.
- Set your LOCK setting to a value that is the largest number required.

Informix ONCONFIG File Recommendations and Best Practices

Advanced DataTools

Basic Informix ONCONFIG Setup and Configuration

ROOTPATH	/PATH/rootdbs
ROOTSIZE	400000
MSGPATH	/PATH/trainX_online.log
CONSOLE	/PATH/trainX_console.log
TAPEDEV	/dev/null
LTAPEDEV	/dev/null
SERVERNUM	<Your Server number goes here>
DBSERVERNAME	<Informix Server Name>
DBSERVERALIASES	<Informix Server Alias Name>
NETTYPE	ipcshm,1,50,CPU
NETTYPE	soctcp,1,50,NET

Basic Informix ONCONFIG Setup and Configuration

- ROOTPATH – full path location to your rootdbs
  - ROOTPATH /informixchunks/train1/rootdbs
- ROOTSIZE – Size of your rootdbs in KB
  - ROOTSIZE 2000000
- File must be owned by Informix and belong to the Informix group
- File Permissions must be read/write by user and group Informix only

Basic Informix ONCONFIG Setup and Configuration

- MSGPATH – Full path to the location of the Informx Message log file
  - MSGPATH $INFORMIXDIR/train1_online.log
- CONSOLE – Full path to the location of the Informx Console log file
  - CONSOLE $INFORMIXDIR/train1_console.log

Basic Informix ONCONFIG Setup and Configuration

- Set Ontape Backups to a directory
  - TAPEDEV /home/informix/backups/servername/archive
  - LTAPEDEV /home/informix/backups/servername/logs
- Directory must be owned by Informix and belong to the Informix group
- Permissions must be read/write by user and group Informix only

Basic Informix ONCONFIG Setup and Configuration

- SERVERNUM – Must be a unique number for each instance on a machine
- DBSERVERNAME – the Server Name
  - The connections INFORMIXSERVER
- DBSERVERALIAS – the Server Alias Name for other (Network Connections)
  - The connections INFORMIXSERVER

Basic Informix ONCONFIG Setup and Configuration

- NETTYPE – The Network settings for your Server – See Slide 26-27

Disk Space Configuration

- Root DBspace
- Physical Log DBspace
- Logical Logs Dbspace
- Temp Dbspace
- Data Dbspace
- Index DBspace

ONCONFIG Setting

- The following are some Critical ONCONFIG File Settings to Review

ONCONFIG - Root DBspace

```
#########################################################################
# Root Dbspace Configuration Parameters
#########################################################################
# ROOTNAME - The root dbspace name to contain reserved pages and internal tracking tables.
# ROOTPATH - The path for the device containing the root dbspace
# ROOTOFFSET - The offset, in KB, of the root dbspace into the device. The offset is
# required for some raw devices.
# ROOTSIZE - The size of the root dbspace, in KB. The value of
# 200000 allows for a default user space of about
# 100 MB and the default system space requirements.
# MIRROR - Enable (1) or disable (0) mirroring
# MIRRORPATH - The path for the device containing the mirrored root dbspace
# MIRROROFFSET - The offset, in KB, into the mirrored device
# Warning: Always verify ROOTPATH before performing disk initialization (oninit -i or -iy)
# to avoid disk corruption of another instance
#########################################################################

ROOTNAME        rootdbs
ROOTPATH         /informixchunks/rootdbs
ROOTOFFSET       0
ROOTSIZE        200000
MIRROR          1
MIRRORPATH      /informixchunks/rootdbsM
MIRROROFFSET    0
```

ONCONFIG - Physical Log

```
#------------------------------------------------------------------
# Physical Log Configuration Parameters
#------------------------------------------------------------------
# PHYSFILE - The size, in KB, of the physical log on disk.
# If RTO_SERVER_RESTART is enabled, the
# suggested formula for the size of PHSYFILE
# (up to about 1 GB) is:
# PHYSFILE = Size of BUFFERS * 1.1
# PLOG_OVERFLOW_PATH - The directory for extra physical log files
# if the physical log overflows during recovery
# or long transaction rollback
# PHYSBUFF - The size of the physical log buffer, in KB
#------------------------------------------------------------------

PHYSFILE              50000
PLOG_OVERFLOW_PATH $INFORMIXDIR/tmp
PHYSBUFF              128
```

ONCONFIG - Logical Log

```
#############################################################
# Logical Log Configuration Parameters
#############################################################
# LOGFILES - The number of logical log files
# LOGSIZE - The size of each logical log, in KB
# DYNAMIC_LOGS - The type of dynamic log allocation.
# Acceptable values are:
# 2 Automatic. IDS adds a new logical log to the
# root dbspace when necessary.
# 1 Manual. IDS notifies the DBA to add new logical
# logs when necessary.
# 0 Disabled
# LOGBUFF - The size of the logical log buffer, in KB
#############################################################

LOGFILES        6
LOGSIZE         10000
DYNAMIC_LOGS    2
LOGBUFF         64
```

ONCONFIG - Long Transactions

```
#############################################################
# Long Transaction Configuration Parameters
#############################################################
# If IDS cannot roll back a long transaction, the server hangs
# until more disk space is available.
# LTXHWM - The percentage of the logical logs that can be filled before a
# transaction is determined to be a long transaction and is rolled back
# LTXEHWM - The percentage of the logical logs that have been filled before the
# server suspends all other transactions so that the long transaction being rolled back
# has exclusive use of the logs
# When dynamic logging is on, you can set higher values for LTXHWM and LTXEHWM because the
# server can add new logical logs during long transaction rollback. Set lower values to
# limit the number of new logical logs added.
#
# If dynamic logging is off, set LTXHWM and LTXEHWM to lower values, such as 50 and 60 or
# lower, to prevent long transaction rollback from hanging the server due to lack of
# logical log space.
#
# When using Enterprise Replication, set LTXEHWM to at least 30%
# higher than LTXHWM to minimize log overruns.
#############################################################
# NOTE: The new default is LTXHWM 70, LTXEHWM 80 – I recommend using the following values
LTXHWM 50
LTXEHWM 60
```

ONCONFIG - Server Name and Aliases

```
##########################################################################
# System Configuration Parameters
##########################################################################
# SERVERNUM - The unique ID for the IDS instance. Acceptable
# values are 0 through 255, inclusive.
# DBSERVERNAME - The name of the default database server
# DBSERVERALIASES - The list of up to 32 alternative dbservernames,
# separated by commas
##########################################################################
SERVERNUM 1
DBSERVERNAME train
DBSERVERALIASES train2_tcp/3, train_drda
```

Following a soctcp connection name with a slash and a number causes the engine to start that number of listener threads for this port. The default is a single listener thread.

ONCONFIG - CPU Configuration

```
##########################################################################
# CPU-Related Configuration Parameters
##########################################################################
# MULTIPROCESSOR - Specifies whether the computer has multiple CPUs. Acceptable
# values are: 0 (single processor), 1 (multiple processors or multi-core chips)
# VPCLASS cpu - Configures the CPU VPs. The format is:
# VPCLASS cpu,num=<#>[,max=<#>][,aff=<#>] [,noage]
# VP_MEMORY_CACHE_KB - Specifies the amount of private memory blocks of your CPU VP, in
# KB, that the database server can access. Acceptable values are:
# 0 (disable)
# 800 through 40% of the value of SHMTOTAL
# SINGLE_CPU_VP - Optimizes performance if IDS runs with only one CPU VP. Acceptable
# values are:
# 0 multiple CPU VPs
# Any nonzero value (optimize for one CPU VP)
##########################################################################

MULTIPROCESSOR 0
VPCLASS cpu,num=1,noage
VP_MEMORY_CACHE_KB 0
SINGLE_CPU_VP 1
```

ONCONFIG - Network Connections

```
#########################################################################
# NETTYPE - The configuration of poll threads
# for a specific protocol. The format is:
# NETTYPE <protocol>,<# poll threads>
# ,<number of connections/thread>,(NETICPU)
# You can include multiple NETTYPE
# entries for multiple protocols.
# LISTEN_TIMEOUT - The number of seconds that IDS
# waits for a connection
# MAX_INCOMPLETE_CONNECTIONS - The maximum number of incomplete
# connections before IDS logs a Denial
# of Service (DoS) error
# FASTPOLL - Enables (1) or disables (0) fast
# polling of your network, if your
# operating system supports it.
# NS_CACHE - The number of seconds for IDS name service cache
# (host, service, user, group) expiration time.
# 0 to disable cache.
#########################################################################
NETTYPE              ipcshm,1,50,CPU
NETTYPE              soctcp,1,50,NET
LISTEN_TIMEOUT       60
MAX_INCOMPLETE_CONNECTIONS 1024
FASTPOLL             1
NS_CACHE             host=900,service=900,user=900,group=900
```

ONCONFIG - Checkpoints and Recover

```
##################################################################
# Checkpoint and System Block Configuration Parameters
##################################################################
# CKPINTVL - Specifies how often, in seconds, IDS checks
# if a checkpoint is needed. 0 indicates that
# IDS does not check for checkpoints. Ignored
# if RTO_SERVER_RESTART is set.
# AUTO_CKPTS - Enables (1) or disables (0) monitoring of
# critical resource to trigger checkpoints
# more frequently if there is a chance that
# transaction blocking might occur.
# RTO_SERVER_RESTART - Specifies, in seconds, the Recovery Time
# Objective for IDS restart after a server
# failure. Acceptable values are 0 (off) and
# any number from 60-1800, inclusive.
# BLOCKTIMEOUT - Specifies the amount of time, in seconds,
# for a system block.
##################################################################
CKPTINTVL 300
AUTO_CKPTS 1
RTO_SERVER_RESTART 0
BLOCKTIMEOUT 3600
```

ONCONFIG – Auto Tune

```
#------------------------------------------------------------------------
# AUTO_TUNE    - The value of this parameter serves as the default value for
#                the following AUTO_* parameters:
#                AUTO_AIOVPS
#                AUTO_CKPTS
#                AUTO_REPREPARE
#                AUTO_STAT_MODE
#                AUTO_READAHEAD
#                AUTO_LRU_TUNING
#
# Any of the above parameters that are not present in your config file
# will default to the value of AUTO_TUNE, which can be set to either 0 or 1.
# If an AUTO_* parameter is set in your config file, the given value overrides
# that of AUTO_TUNE. Information on individual AUTO_* parameters is below.
#
```

ONCONFIG – Auto Tune

```
# AUTO_LRU_TUNING - Enables (1) or disables (0) automatic LRU tuning, which
#                   adjusts flushing thresholds for individual buffer pools
#                   if the server discovers they are sub-optimal
# AUTO_AIOVPS     - Enables (1) or disables (0) automatic management
#                   of AIO VPs# AUTO_CKPTS      - Enables (1) or disables (0) monitoring of
#                   critical resource to trigger checkpoints
#                   more frequently if there is a chance that
#                   transaction blocking might occur.
# AUTO_REPREPARE  - Enables (1) or disables (0) automatically
#                   re-optimizing stored procedures and re-preparing
#                   prepared statements when tables that are referenced
#                   by them change. Minimizes the occurrence of the
#                   -710 error.# AUTO_STAT_MODE  - Enables (1) or disables (0) update statistics
#                   automatic mode. In automatic mode, statistics of
#                   table, fragment or index are rebuilt only if existing
#                   statistics are considered stale. A table, fragment
#                   or index can change by STATCHANGE percentage before
#                   its statistics are regarded as stale.
```

ONCONFIG – Auto Tune

```
# RA_PAGES & RA_THRESHOLD have been replaced with AUTO_READAHEAD.
# AUTO_READAHEAD mode[,readahead_cnt]
#     mode         0 = Disable     (Not recommended)
#                  1 = Passive     (Default)
#                  2 = Aggressive  (Not recommended)
#     readahead_cnt     Optional        Range 4-4096
#                  readahead_cnt allows for tuning the of
#                  pages that automatic readahead will request
#                  to be read ahead. When not set, the default
#                  is 128 pages.
#     Notes:
#     The threshold for starting the next readahead request, which
#     used to be known as RA_THRESHOLD, is always set to 1/2 of the
#     readahead_cnt. RA_THRESHOLD is deprecated and no longer used.
# If RA_PAGES & AUTO_READAHEAD are not present in the ONCONFIG file,
# AUTO_READAHEAD will default to the value of AUTO_TUNE.
# If RA_PAGES is present in the ONCONFIG file and AUTO_READAHEAD is
# not, the server will set AUTO_READAHEAD to AUTO_TUNE,RA_PAGES
```

ONCONFIG - Transactions

```
#########################################################################
# Transaction-Related Configuration Parameters
#########################################################################
# TXTIMEOUT - The distributed transaction timeout, in seconds
# DEADLOCK_TIMEOUT - The maximum time, in seconds, to wait for a
# lock in a distributed transaction.
# HETERO_COMMIT - Enables (1) or disables (0) heterogeneous
# commits for a distributed transaction
# involving an EGM gateway.
#########################################################################
TXTIMEOUT 300
DEADLOCK_TIMEOUT 60
HETERO_COMMIT 0
```

ONCONFIG - Disk I/O

```
##############################################################
# AIO and Cleaner-Related Configuration Parameters
##############################################################
# VPCLASS aio - Configures the AIO VPs. The format is:
# VPCLASS aio,num=<#>[,max=<#>][,aff=<#>][,noage]
# CLEANERS - The number of page cleaner threads
# AUTO_AIOVPS - Enables (1) or disables (0) automatic management of AIO VPs
# DIRECT_IO - Specifies whether direct I/O is used for cooked files used for dbspace chunks.
# Acceptable values are:
# 0 Disable
# 1 Enable direct I/O
# 2 Enable concurrent I/O – AIX Only
##############################################################
#VPCLASS aio,num=1
CLEANERS 8
AUTO_AIOVPS 1
DIRECT_IO 1
```

ONCONFIG - Table Space

```
#*****************************************************************
# Tblspace Configuration Parameters
#*****************************************************************
# TBLTBLFIRST - The first extent size, in KB, for the tblspace
# tblspace. Must be in multiples of the page size.
# TBLTBLNEXT - The next extent size, in KB, for the tblspace
# tblspace. Must be in multiples of the page size.
# The default setting for both is 0, which allows IDS to manage
# extent sizes automatically.
#
# TBLSPACE_STATS - Enables (1) or disables (0) IDS to maintain
# tblspace statistics
#*****************************************************************

TBLTBLFIRST 0
TBLTBLNEXT 0
TBLSPACE_STATS 1
```

ONCONFIG - Temp Space

```
#########################################################################
# Temporary dbspace and sbspace Configuration Parameters
#########################################################################
# DBSPACETEMP - The list of dbspaces used to store temporary
# tables and other objects. Specify a colon
# separated list of dbspaces that exist when the
# server is started. If no dbspaces are specified,
# or if all specified dbspaces are not valid,
# temporary files are created in the /tmp directory
# instead.
# SBSPACETEMP - The list of sbspaces used to store temporary
# tables for smart large objects. If no sbspace
# is specified, temporary files are created in
# a standard sbspace.
#########################################################################

DBSPACETEMP tmp1dbs:tmp2dbs:tmp3dbs:non_temp_space1:non_temp_space2
SBSPACETEMP tmp1sbdbs:tmp2sbdbs:tmp3sbdbs
TEMPTAB_NOLOG  1
```

ONCONFIG - Sbspace and DbSpace

```
##########################################################################
# Dbspace and sbspace Configuration Parameters
##########################################################################
# SBSPACENAME - The default sbspace name where smart large objects
# are stored if no sbspace is specified during
# smart large object creation. Some DataBlade
# modules store smart large objects in this
# location.
# SYSSBSPACENAME - The default sbspace for system statistics
# collection. Otherwise, IDS stores statistics
# in the sysdistrib system catalog table.
# ONDBSPACEDOWN - Specifies how IDS behaves when it encounters a
# dbspace that is offline. Acceptable values
# are:
# 0 Continue
# 1 Stop
# 2 Wait for DBA action
##########################################################################
SBSPACENAME sbspacedbs
SYSSBSPACENAME dbspacedbs
ONDBSPACEDOWN 2
```

ONCONFIG - PDQ Setup

```
##########################################################################
# Parallel Database Query (PDQ) Configuration Parameters
##########################################################################
# MAX_PDQPRIORITY - The maximum amount of resources, as a percentage, that PDQ can
# allocate to any one decision support query
# DS_MAX_QUERIES - The maximum number of concurrent decision support queries
# DS_TOTAL_MEMORY - The maximum amount, in KB, of decision support query memory
# DS_MAX_SCANS - The maximum number of concurrent decision support scans
# DS_NONPDQ_QUERY_MEM - The amount of non-PDQ query memory, in KB. Acceptable values are
# 128 to 25% of DS_TOTAL_MEMORY.
# DATASKIP - Specifies whether to skip dbspaces when
# processing a query. Acceptable values are:
# - ALL Skip all unavailable fragments
# - ON <dbspace1> <dbspace2>... Skip listed
# dbspaces
# - OFF Do not skip dbspaces (default)
##########################################################################
MAX_PDQPRIORITY 100
DS_MAX_QUERIES
DS_TOTAL_MEMORY
DS_MAX_SCANS 1048576
DS_NONPDQ_QUERY_MEM 128
DATASKIP
```

ONCONFIG - Optimizer Setup

```
##########################################################################
# Optimizer Configuration Parameters
##########################################################################
# OPTCOMPIND - Controls how the optimizer determines the best
# query path. Acceptable values are:
# 0 Nested loop joins are preferred
# 1 If isolation level is repeatable read,
# works the same as 0, otherwise works same as 2
# 2 Optimizer decisions are based on cost only
# DIRECTIVES - Specifies whether optimizer directives are
# enabled (1) or disabled (0). Default is 1.
# EXT_DIRECTIVES - Controls the use of external SQL directives.
# Acceptable values are:
# 0 Disabled
# 1 Enabled if the IFX_EXTDIRECTIVES environment
# variable is enabled
# 2 Enabled even if the IFX_EXTDIRECTIVES
# environment is not set
```

Optimizer Setup - continued

```
# OPT_GOAL - Controls how the optimizer should optimize for
# fastest retrieval. Acceptable values are:
# -1 All rows in a query
# 0 The first rows in a query# IFX_FOLDVIEW - Enables (1) or disables (0) folding views that
# have multiple tables or a UNION ALL clause.
# Disabled by default.
# AUTO_REPREPARE - Enables (1) or disables (0) automatically
# re-optimizing stored procedures and re-preparing
# prepared statements when tables that are referenced
# by them change. Minimizes the occurrence of the
# -710 error.
#####################################################################
OPTCOMPIND 2
DIRECTIVES 1
EXT_DIRECTIVES 0
OPT_GOAL -1
IFX_FOLDVIEW 0
AUTO_REPREPARE 1
```

ONCONFIG - Dump Setup

```
#########################################################################
# Diagnostic Dump Configuration Parameters
#########################################################################
# DUMPDIR - The location Assertion Failure (AF) diagnostic
# files
# DUMPSHMEM - Controls shared memory dumps. Acceptable values
# are:
# 0 Disabled
# 1 Dump all shared memory
# 2 Exclude the buffer pool from the dump
# DUMPGCORE - Enables (1) or disables (0) whether IDS dumps a
# core using gcore
# DUMPCORE - Enables (1) or disables (0) whether IDS dumps a
# core after an AF
# DUMPCNT - The maximum number of shared memory dumps or
# core files for a single session
#########################################################################
DUMPDIR $INFORMIXDIR/tmp
DUMPSHMEM 1
DUMPGCORE 0
DUMPCORE 0
DUMPCNT 1
```

ONCONFIG - Alarm Program Setup

```
#############################################################
# Alarm Program Configuration Parameters
#############################################################
# ALARMPROGRAM - Specifies the alarm program to display event
# alarms. To enable automatic logical log backup,
# edit alarmprogram.sh and set BACKUPLOGS=Y.
# ALRM_ALL_EVENTS - Controls whether the alarm program runs for
# every event. Acceptable values are:
# 0 Logs only noteworthy events
# 1 Logs all events
# STORAGE_FULL_ALARM - <time interval in seconds>,<alarm severity>
# specifies in what interval:
# - a message will be printed to the online.log file
# - an alarm will be raised
# when
# - a dbspace becomes full
# (ISAM error -131)
# - a partition runs out of pages or extents
# (ISAM error -136)
# time interval = 0 : OFF
# severity = 0 : no alarm, only message
# SYSALARMPROGRAM - Specifies the system alarm program triggered
# when an AF occurs
#############################################################
```

Alarm Program Setup

ALARMPROGRAM $INFORMIXDIR/etc/log_full.sh
ALRM_ALL_EVENTS 1
STORAGE_FULL_ALARM 600,3
SYSALARMPROGRAM $INFORMIXDIR/etc/evidence.sh

Best Practices Summary

Next Webcast
Informix Best Practices

- Disks and Database Space Layout
  - Thursday, March 30, 2017 at 2:00pm EST
- Backup, Recovery, and High Availability Disaster Recovery
  - Thursday, April 20, 2017 at 2:00pm EST

Please register for each webcast here at:
http://advancedatatools.com/Informix/NextWebcast.html

Informix Resources - IIUG

- The International Informix User Group
  - http://www.iiug.org
  - Membership is FREE
- IIUG 2017 – The Premier Informix Event
  April 23 – 27, 2017
  - http://www.iiug2017.org

Informix Resources from IBM

- Informix Documentation
  http://www.ibm.com/support/knowledgecenter/SSGU8G_12.1.0/com.ibm.welcome.doc/welcome.htm
- Compare the Informix Version 12 editions by Carlton Doe, IBM
  http://www.ibm.com/developerworks/data/library/techarticle/dm-0801doe/
- The Informix and IoT Roadshows by Carlton Doe, IBM
  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/North%20America%20Informix%20Events

Informix Training in 2017

- April 10-13, 2017
  - **Informix for Database Administrators**
- July 10-13, 2017
  - **Advanced Informix Performance Tuning**
- September 18-21, 2017
  - **Informix for Database Administrators**
- All courses can be taken online on the web from your desk or at our training center in Virginia.
- We guarantee to *NEVER* cancel a course and will teach a course as long as one student is registered!

Questions?

Send follow-up questions to
lester@advancedatatools.com

Advanced DataTools

Informix Support and Training from the Informix Champions!

Advanced DataTools is an Advanced Level IBM Informix Data Management Partner, and has been an authorized Informix partner since 1993. We have a long-term relationship with IBM, we have priority access to high-level support staff, technical information, and Beta programs. Our team has been working with Informix since its inception, and includes 8 Senior Informix Database Consultants, 4 IBM Champions, 2 IIUG Director's Award winners, and an IBM Gold Consultant. We have Informix specialists Lester Knutsen and Art Kagel available to support your Informix performance tuning and monitoring requirements!

- ***Informix Remote DBA Support Monitoring***
- ***Informix Performance Tuning***
- ***Informix Training***
- ***Informix Consulting***
- ***Informix Development***

Free Informix Performance Tuning Webcast replays at:
http://advancedatatools.com/Informix/Webcasts.html

Call: (800) 807-6732 x101 or Email: info@advancedatatools.com
Web: http://www.advancedatatools.com

Informix Best Practices

Advanced DataTools

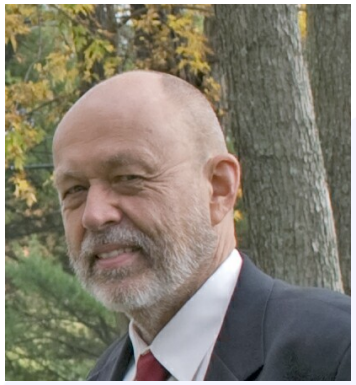

Thank You

Lester Knutsen

Advanced DataTools Corporation

lester@advancedatatools.com

For more information:

http://www.advancedatatools.com

Advanced DataTools