Informix Performance Tuning
Exploring the Sysmaster Database (New)
by Lester Knutsen

Webcast on March 24, 2015 at 2:00pm EST

Lester Knutsen

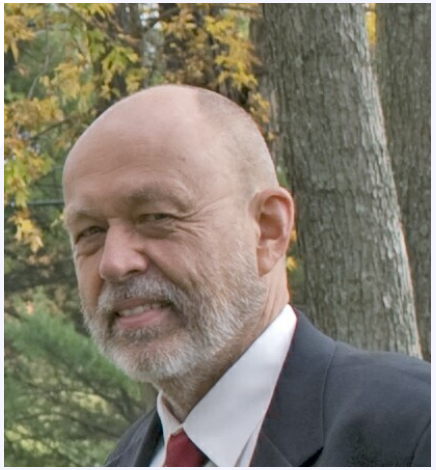

Lester Knutsen is President of Advanced DataTools Corporation, and has been building large Data Warehouse and Business Systems using Informix Database software since 1983. Lester focuses on large database performance tuning, training and consulting. Lester is a member of the IBM Gold Consultant program 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

Agenda

- What is the Sysmaster Database?
- Demo – New Scripts
- New Sysmaster Tables
- Monitoring SQL with SQL Trace

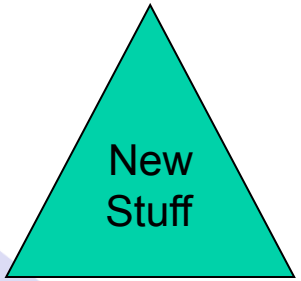

Past Presentations

- Informix Performance Tuning using the Sysmaster Database by Lester Knutsen
  - Tuesday, April 30, 2013
  - 101 pages of slides and scripts
- Webcast Replay on YouTube:
  - http://youtu.be/eFr5cP2Ly0s
  - http://advancedatatools.com/Informix/Webcasts.html

New Sysmaster Queries

- What is the Hardware?
- What Informix Features are used?
- What is the SQLHosts file?
- What is the Informix Memory Usage?
- What is the Informix Oninit CPU Usage?
- What is the Disk IO and History?
- What is the Checkpoint History?

New Sysmaster Queries

- What is the table layout on disk?
- What tables need to be reorganized due to fragmentation and extents?
- What tables have sequential Scans?
- Where are the Logical Logs?
- What threads are running?
- What are the most costly queries?

What is the Sysmaster Database?

A database that peeks into
the shared memory structures
of an INFORMIX-Dynamic Server

Informix Control Structures in Memory are the Sysmaster Database

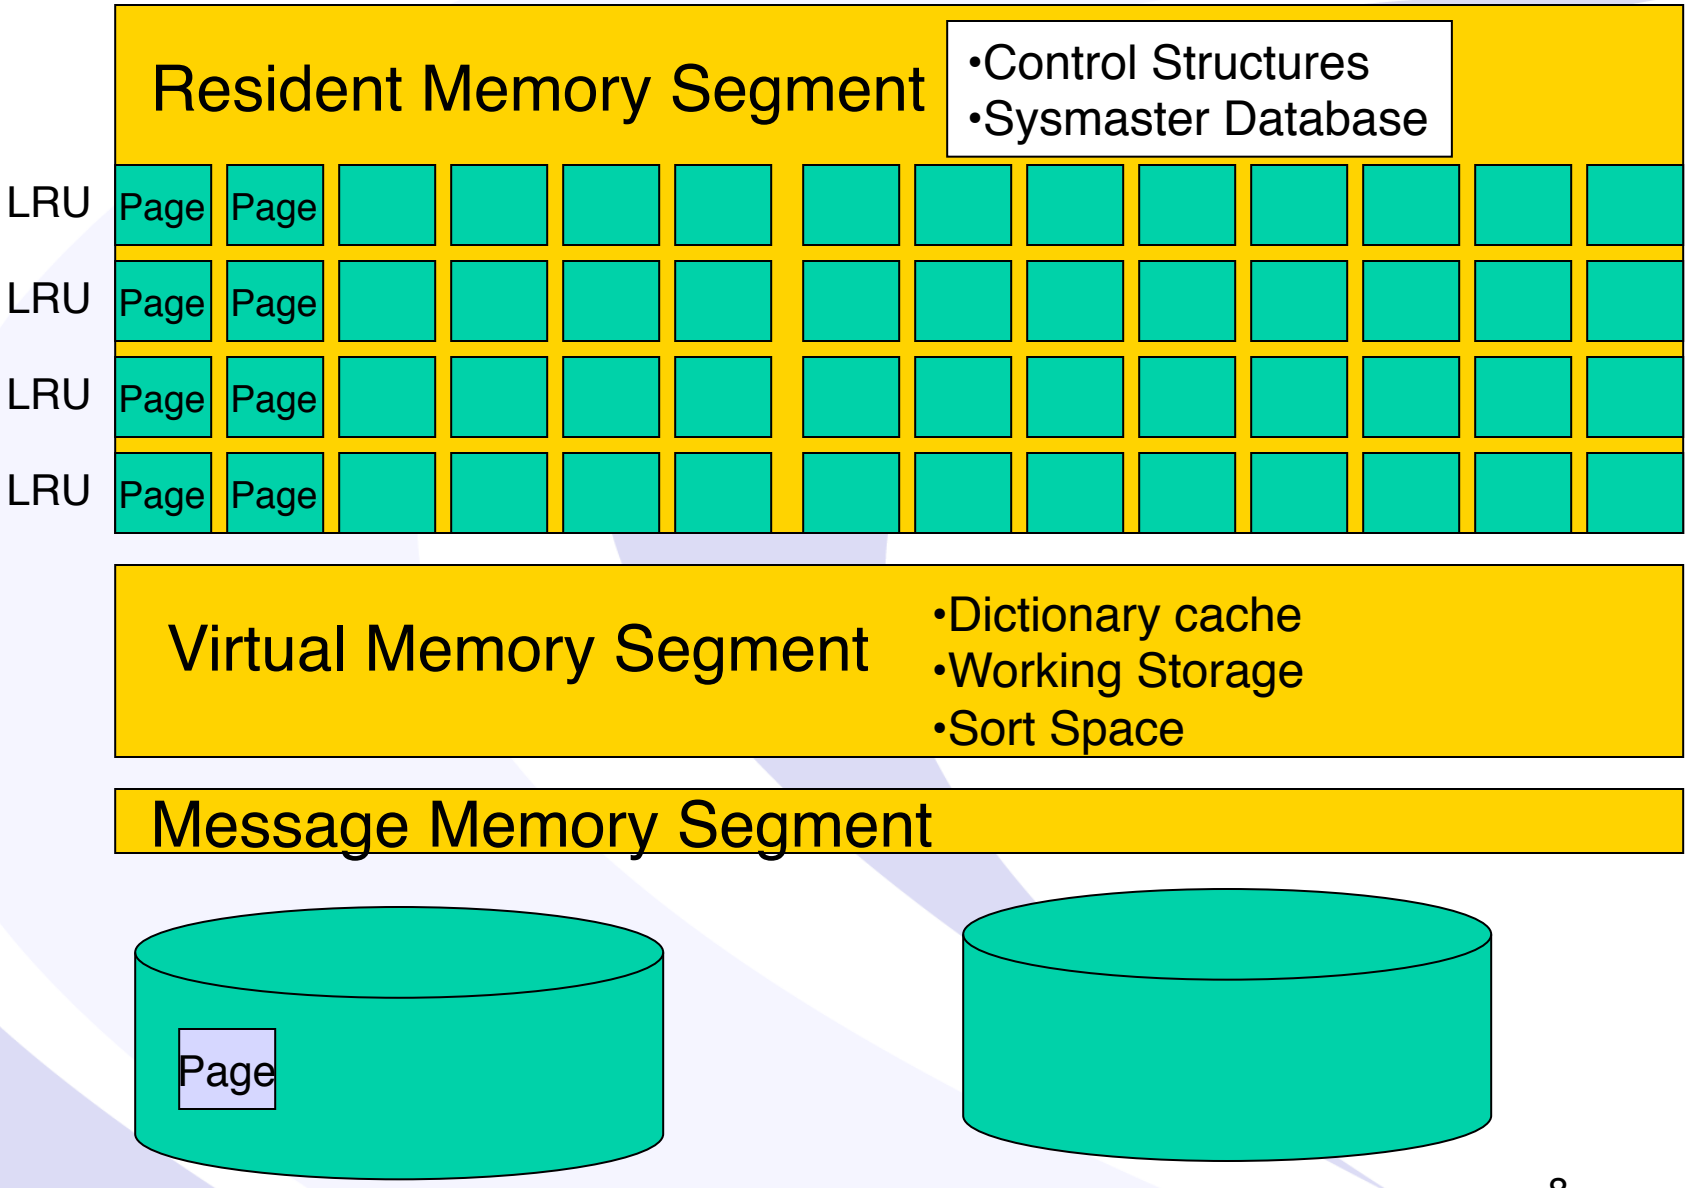

Sysmaster Database contains:

- Server information
- Dbspace & chunk information
- Database & table information
- User session information

Performance of queries on Sysmaster Database

The data is in shared memory but:

- Views used by tables require disk access and may be slow
- Complex views used to hide complex data
- Some tables are large (million locks)
- Unbuffered logging of temp tables

Differences from other databases

- Do not update Sysmaster tables as this may corrupt the server
- Cannot use dbschema on pseudo tables
- Cannot drop pseudo tables or the Sysmaster Database

Isolation level is Dirty Read

- Data is dynamic and can change as you retrieve it (Dirty Read)
- Dynamic nature may return inconsistent results
- However, it uses Unbuffered logging and temp tables are logged

Sysmaster Database may change

- Some undocumented tables and columns may change in future versions
- Scripts in this presentation using undocumented features may not work on all versions of Informix
- New scripts have been run on versions 11.7 and 12.X
- Sysmaster has changed in IDS 11.X and 12.X

New Scripts and Tables

What is the Hardware? Table:Sysmachineinfo*

```
os_name             char(128),
os_release          char(128),
os_nodename         char(128),
os_version          char(128),
os_machine          char(128),
os_num_procs        smallint,
os_num_olprocs      smallint,
os_pagesize         int8,
os_mem_total        int8,
os_mem_free         int8,
os_open_file_lim    int8,
os_shmmax           int8,
os_shmmin           int8,
os_shmids           int8,
os_shmnumsegs       int8,
os_semmap           int8,
os_semids           int8,
os_semnum           int8,
os_semundo          int8,
os_semnumperid      int8,
os_semops           int8,
os_semundoperproc   int8,
os_semundosize      int8,
os_semmaxvalue      int8
```

What is the Hardware? Table:Sysmachineinfo *

os_name	Linux
os_release	3.16.7-7-desktop
os_nodename	train6
os_version	#1 SMP PREEMPT Wed Dec 17 18:00:44 UTC 2014 (762f27a)
os_machine	x86_64
os_num_procs	4
os_num_olprocs	4
os_pagesize	4096
os_mem_total	6271201280
os_mem_free	5420998656
os_open_file_lim	32768
os_shmmax	4398046511104
os_shmmin	1
os_shmids	4096
os_shmnumsegs	4194304
os_semmap	
os_semids	128
os_semnum	128000
os_semundo	
os_semnumperid	250
os_semops	100
os_semundoperproc	
os_semundosize	20
os_semmaxvalue	32767

What Informix Features are used? Table: Syslicenseinfo*

```
version           char(12),    --  Informix version
week              smallint,    --  Week
year              smallint,    --  Year
max_cpu_vps       smallint,    --  Max number of cpu vps
max_vps           smallint,    --  Max number of vps
max_conns         integer,     --  Max # of user connected
max_sec_conns     integer,     --  Max # of secondary user
max_sds_conns     integer,     --  Max # of sds users
max_sds_clones    smallint,    --  Max # of sds clones
max_rss_clones    smallint,    --  Max # of rss clones
total_size        integer,     --  max disk space (MB)
total_size_used   integer,     --  max disk space used (MB)
max_memory        integer,     --  Max memory allocated (MB)
max_memory_used   integer,     --  Max memory used (MB)
feature_flags     integer,     --  Feature Flags
feature_flags2    integer      --  Feature Flags2
```

What Informix Features are used? View: Sysfeatures

```
create view sysfeatures (                  { Internal Use Only              }
          week, year, version, max_cpu_vps, max_vps,
          max_conns, max_sec_conns, max_sds_clones, max_rss_clones,
          total_size, total_size_used,
          max_memory, max_memory_used, is_primary, is_secondary,
          is_sds, is_rss, is_er, is_pdq )
   AS
   select  week, year, version, max_cpu_vps, max_vps,
          max_conns, max_sec_conns, max_sds_clones, max_rss_clones,
          format_units(total_size,'M'),
          format_units(total_size_used,'M'),
          format_units(max_memory,'M'),
          format_units(max_memory_used,'M'),
          decode(bitand(feature_flags, 1),0,0,1),
          decode(bitand(feature_flags, 2),0,0,1),
          decode(bitand(feature_flags, 8),0,0,1),
          decode(bitand(feature_flags, 4),0,0,1),
          decode(bitand(feature_flags, 16),0,0,1),
          decode(bitand(feature_flags, 512),0,0,1)
   from syslicenseinfo
```

What Informix Features are used? View: Sysfeatures

```
week             12
year             2015
version          12.10.FC4
max_cpu_vps      1
max_vps          28
max_conns        2
max_sec_conns    0
max_sds_clones   0
max_rss_clones   0
total_size       26.6 GB
total_size_used  7.16 GB
max_memory       1.27 GB
max_memory_used  1.25 GB
is_primary       0
is_secondary     0
is_sds           0
is_rss           0
is_er            0
is_pdq           0
```

What Informix Features are used? View: Sysfeatures

year	week	version	max_cpu_vps	max_conns	max_memory
2015	12	12.10.FC4	4	8	5417
2015	11	12.10.FC4	4	9	5417
2015	10	11.70.UC3	4	6	2837
2015	9	11.70.UC3	4	5	2837
2015	8	11.70.UC3	4	7	2837
2015	7	11.70.UC3	4	7	2837
2015	6	11.70.UC3	4	2	2837
2015	5	11.70.UC3	4	2	2837
2015	4	11.70.UC3	4	2	2837
2015	3	11.70.UC3	4	2	2837
2015	2	11.70.UC3	4	3	2837
2015	1	11.70.UC3	4	2	2837
2014	52	11.70.UC3	4	1	2837
2014	51	11.70.UC3	4	4	3032
2014	50	11.70.UC3	4	6	3032
2014	49	11.70.UC3	4	7	3032
2014	48	11.70.UC3	4	6	3032

What is the SQLHosts file? Table: Syssqlhosts

```
dbsvrnm             char(129),
nettype             char(17),
svrtype             char(9),
netprot             char(9),
hostname            char(257),
svcname             char(129),
options             char(257),
svrsecurity         smallint,
clntsecurity        smallint,
netoptions          smallint,
netbuf_size         integer,
connmux_option      smallint,
svrgroup            char(129),
endofgroup          char(129),
redirector          smallint,
svrid               integer,
pamauth             integer,
authtoken           char(256)
```

What is the SQLHosts file?
Table: Syssqlhosts

```
dbservername  train1
nettype       onipcshm
hostname      train6
servicename   train1

dbservername  train1tcp
nettype       onsoctcp
hostname      train6
servicename   1527

dbservername  train6
nettype       onipcshm
hostname      train6
servicename   train6

dbservername  train6tcp
nettype       onsoctcp
hostname      train6
servicename   sqlexec
```

What is the Memory Usage? Table:Sysseglst

```
seg_address         int8,     --  address of segment structure
seg_next            int8,     --  pointer to next segment
seg_prev            int8,     --  pointer to prev segment
seg_class           smallint, --  segment class
seg_size            int8,     --  size of this segment
seg_osshmid         integer,  --  id of this OS segment in this seg
seg_osmaxsize       int8,     --  size of maximum OS segment in this seg
seg_osshmkey        integer,  --  shmkey for first OS segment
seg_procid          integer,  --  process id of creator
seg_userid          smallint, --  usr id of creator
seg_shmaddr         int8,     --  address of segment
seg_ovhd            int8,     --  amount of overhead bytes
seg_lock            int8,     --  lock to synchronise bitmap access
seg_nextid          integer,  --  segment id of next seg
seg_bmapsz          int8,     --  size of block map
seg_blkused         int8,     --  no. of used blocks in segment
seg_blkfree         int8      --  no. of free blocks in segment
```

What is the Memory Usage? Table:Sysseglst

```
-- Summary by Memory Segments Class

select
        -- seg_class,
        case
                when seg_class = 1 then "Resident"
                when seg_class = 2 then "Virtual"
                when seg_class = 3 then "Message"
                when seg_class = 4 then "Buffer"
                else "Unknown"
        end class,
        count(*) number ,
        sum( seg_size ) total_size,
        sum( seg_blkused )      total_blkused,
        sum( seg_blkfree )      total_blkfree
from sysseglst
group by 1;
```

What is the Memory Usage? Table:Sysseglst

```
-- Detail by Memory Segment

select
        -- seg_class,
        case
                when seg_class = 1 then "Resident"
                when seg_class = 2 then "Virtual"
                when seg_class = 3 then "Message"
                when seg_class = 4 then "Buffer"
                else "Unknown"
        end class,
        seg_size,
        seg_blkused,
        seg_blkfree
from sysseglst;
```

What is the Memory Usage? Table:Sysseglst

class	number	total_size	total_blkused	total_blkfree
Message	1	561152	136	1
Resident	1	4902912	1197	0
Buffer	2	1291595776	315331	0
Virtual	5	66994176	12241	4115

class	seg_size	seg_blkused	seg_blkfree
Resident	4902912	1197	0
Virtual	33439744	8160	4
Buffer	1125236736	274716	0
Buffer	166359040	40615	0
Message	561152	136	1
Virtual	8388608	2006	42
Virtual	8388608	1795	253
Virtual	8388608	228	1820
Virtual	8388608	25	2023

What is the Informix Oninit CPU Usage? Table:Sysvplst*

```
pid                 integer,    --  VP id
address             int8,       --  address of VP struct
pid                 integer,    --  unix process id
usecs_user          float,      --  number of usecs of user time
usecs_sys           float,      --  number of usecs of system time
scputimep           int8,       --  ptr to saved cputime (tms)
rcputimep           int8,       --  ptr to reset cputime (tms)
class               integer,    --  class of VP
classname           char(19),   --  classname of VP
readyqueue          int8,       --  ptr to ready queue tab (TCB_Q)
num_ready           integer,    --  number of ready threads
flags               integer,    --  VP flags
next                int8,       --  next in idle list
prev                int8,       --  prev in idle list
semid               integer,    --  semid for this VP
lock                int8,       --  VP protection
total_semops        int8,       --  Total times VP slept on a semop
total_busy_wts      int8,       --  Total VP busy waits
total_yields        int8,       --  Total VP yields
total_spins         int8,       --  Total spins while busy waiting
```

What is the Informix Oninit CPU Usage? Table:Sysvplst*

```
Continued:
steal_attempts          int8,                   --
steal_attempts_suc      int8,                   --
idle_search             int8,                   --
idle_search_suc         int8,                   --
vp_poll_scheds          int8,                   --
vp_mt_naps              int8,                   --
vp_cache_size           int8,                   --  size of the vp cache
vp_cache_allocs         int8,                   --
vp_cache_miss           int8,                   --
vp_cache_frees          int8,                   --
vp_cache_drain          int8,                   --
vp_cache_nblocks        int8,                   --  current number of blocks
thread_run              float,                  --  total thread run time on vp
thread_idle             float,                  --  total time running idle thread
thread_poll_idle        float                   --  inline poll thread idle time
```

What is the Disk I/O and History? Table: Sysiohistory

```
address              bigint,
gfd                  int,
iskaio               int,
open_mode            int,
open_time            bigint,
path                 char(256),
minute               int,
time                 bigint,
total_read_ops       bigint,
total_read_time      float,
read_ops_minute      bigint,
read_time_minute     float,
avg_read_time_minute  float,
total_write_ops      bigint,
total_write_time     float,
write_ops_minute     bigint,
write_time_minute    float,
avg_write_time_minute  float,
total_lseek_time     float,
lseek_time_minute    float
```

Contains last hour of I/O history

What is the Checkpoint History? Table:Syscheckpoint

```
intvl               int,      -- checkpoint interval
type                char(12), -- checkpoint type
caller              char(10), -- caller
clock_time          int,      -- time of day of ckpt
crit_time           float,    -- time spent in wait4critex
flush_time          float,    -- time spent flushing pages to disk
cp_time             float,    -- time from cpkt_pending to done
n_dirty_buffs       int,      -- number of dirty buffers
plogs_per_sec       int,      -- avg # pages plogged
llogs_per_sec       int,      -- avg # pages logged
dskflush_per_sec    int,      -- avg # pages dskflushed
ckpt_logid          int,      -- LSN of ckpt
ckpt_logpos         int,      -- LSN of ckpt
physused            int,      -- total pages plogged in ckpt
logused             int,      -- total pages llogged in ckpt
n_crit_wait         int,      --  # of crit section waiters
tot_crit_wait       float,    -- total time spent waiting for crit
longest_crit_wait   float,    -- longest crit wait
block_time          float     -- blocked time
```

What is the table layout on disk?

```
select    dbinfo ("DBSPACE", pe_partnum ) dbspace,
          pe_chunk  chunknum,
          pe_offset ext_start,
          dbsname   database,
          tabname   partname,
          pe_partnum          partnum,
          pe_extnum extnum,
          pe_size             ext_size,
          pe_log,
          (pe_log + pe_size ) nextone
from      sysptnext b, outer systabnames a
where     a.partnum = b.pe_partnum
order by 1, 2, 3
```

What tables need to be reorganized due to extents?

```
select ( dbinfo('dbspace', ti_partnum )) dbspace,
        dbsname database,
        owner,
        tabname,
        ti_partnum      partnum,
        ti_pagesize     pagesize,
        ti_nptotal      total_pages,
        ti_npused       used_pages,
        ti_npdata       data_pages,
        ti_nextns       num_extents
from systabnames, systabinfo
where ti_partnum = partnum
order by 10 desc;
```

What tables have sequential scans?

```
select  dbsname database,
        tabname,
        ( dbinfo('dbspace', ti_partnum )) dbspace,
        ti_npdata       pages_data,
        seqscans        num_seqscans,
        ti_npdata * seqscans   pages_scanned
from    sysptprof, systabinfo
where   partnum = ti_partnum
and     seqscans > 0
order   by 6 desc;
```

Where are the Logical Logs?

```
select    name dbspace,
          chunk chunknum,
          hex(address) address,
          a.number,
          a.uniqid,
          a.offset,
          a.size,
          a.used,
          a.flags,
          bitval(a.flags, '0x1') used,
          bitval(a.flags, '0x2')  current,
          bitval(a.flags, '0x4')  backedup,
          bitval(a.flags, '0x8')  new,
          bitval(a.flags, '0x10') archived,
          bitval(a.flags, '0x20') temp,
          bitval(a.flags, '0x40') dropped,
          DBINFO ('utc_to_datetime', filltime ) timefull
from syslogfil a, syschunks c, sysdbspaces d
where a.chunk = c.chknum
and c.dbsnum = d.dbsnum
```

What threads are running? Table: sysrstcb

- RSAM Thread Control Block
- Everything you want to know about all running threads….

- Select * from sysrstcb

What threads are running? Table: sysrstcb

Some of the fields in Sysrstcb

```
uid                integer,  -- user id
username           char(32), -- user name
sid                integer,  -- session id
tid                integer,  -- thread id
lkwait             int8,     -- waiting for this lock
lkwttype           integer,  -- lock type waiting for
bfwait             int8,     -- waiting for this buffer
bfwtflag           smallint, -- buffer wait type flag
txwait             int8,     -- waiting for this transaction
txsusp             int8,     -- suspended transaction
nreads             integer,  -- number of reads
nwrites            integer,  -- number of writes
nlocks             integer,  -- number of locks currently held
lkwaittime         float,    -- time spent waiting on locks
iowaittime         float,    -- time spent waiting on disk io
upf_niowaits       integer,  -- Number of disk IO waits
upf_idxbufreads    integer   -- Number of index buffer reads
```

What are the most costly queries? View Sqexplain

- Documented View
- Based on the undocumented table syssdblock

What are the most costly queries? View Sqexplain

```
select
        sqx_estcost,
        sqx_sqlstatement
from    syssqexplain
into    temp A;

select
        sqx_sqlstatement sqlstatement,
        sum(sqx_estcost) sum_estcost,
        count(*)         count_executions
from A
group by 1
order by 2 desc;
```

Informix Performance Tuning using SQL Trace

SQL Trace

- New Feature in Informix 11 to collect SQL statements,
  - Statistics
  - resource usage
  - performance measures
- Default is OFF

Informix Memory Structures

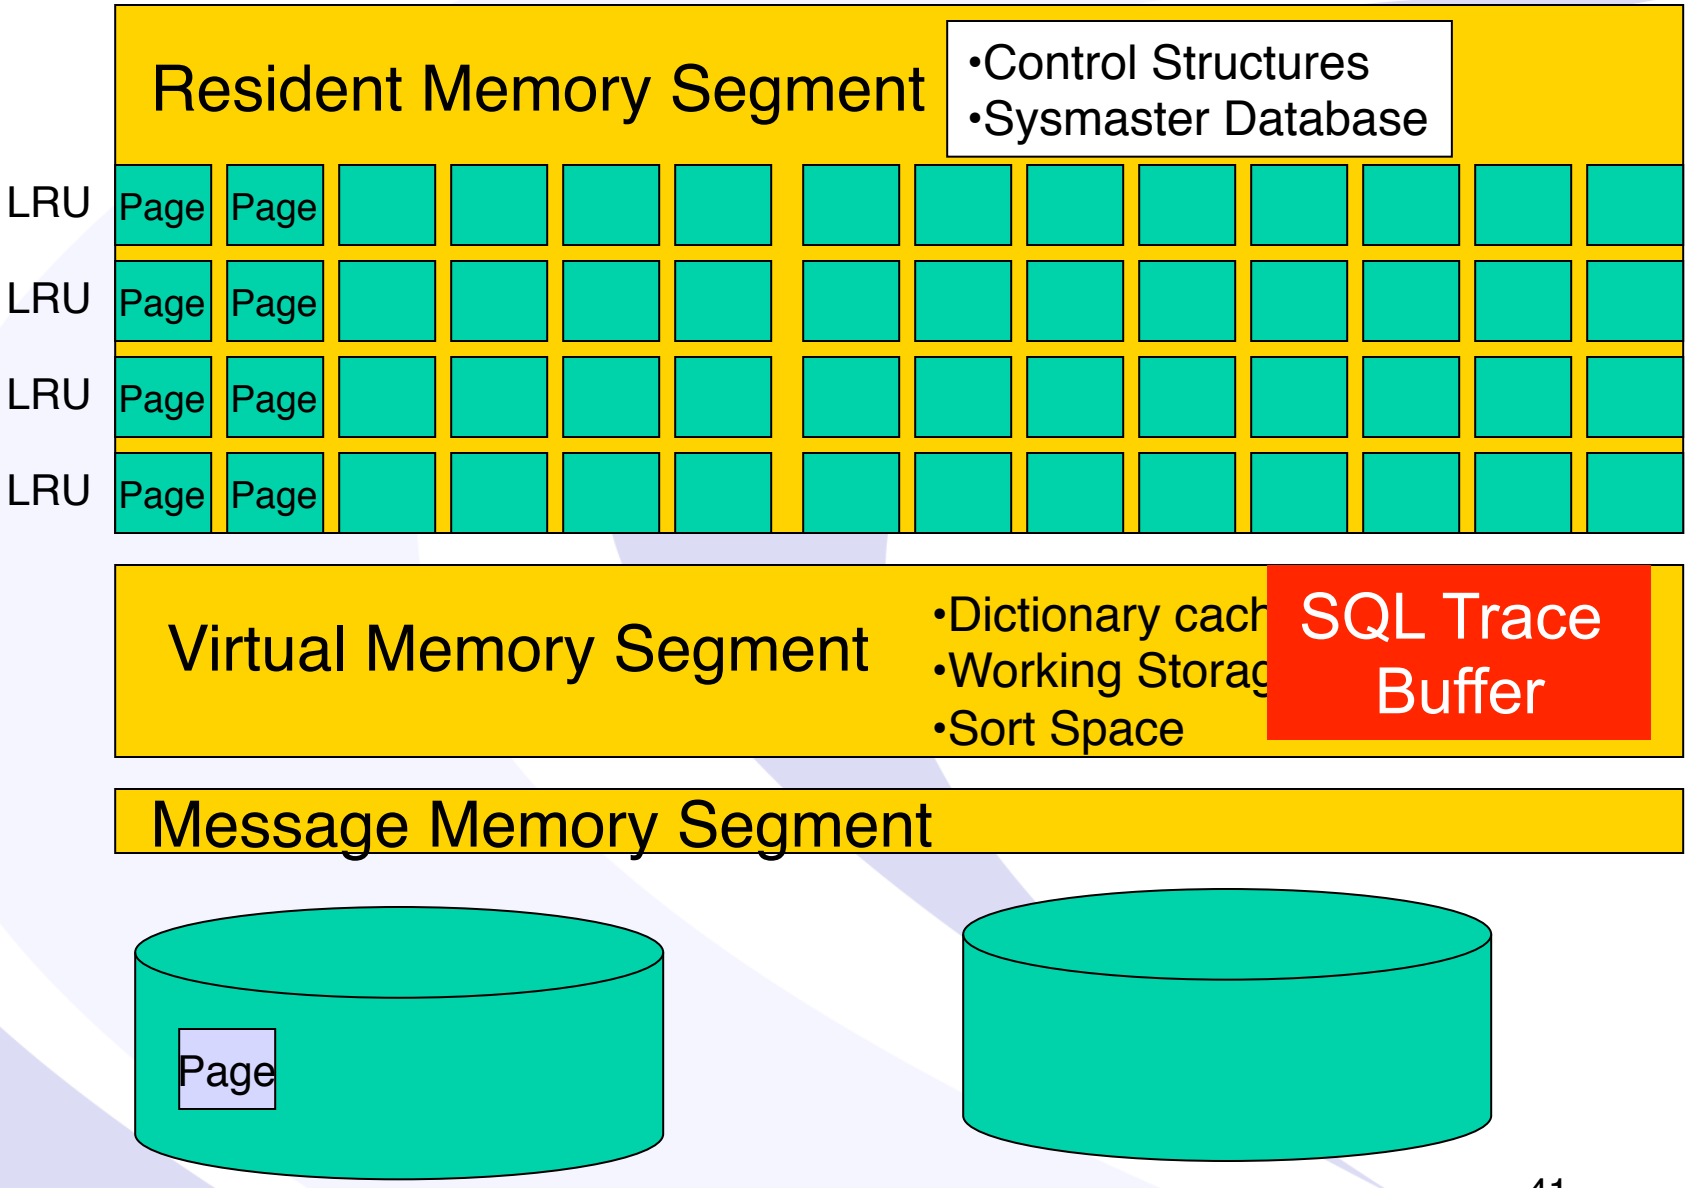

SQL Trace Buffer

- When SQL Trace is turned "ON" an FIFO buffer is created in the Virtual Segment of Informix Memory
- Oldest SQL data is discarded to make room for new data
- On a busy system, the buffer can fill up and turnover in seconds
- Default amount of memory is 2 MB

SQL Trace Configuration

Level - Amount of data to collect

- Off – Default - No SQL tracing
- Low – Default when On. Collects statement statistics, statement text, and statement iterators
- Medium - Collects all low-level tracing, plus table names, the database name, and stored procedure stacks
- High - Collects all of the information included in medium tracing, plus host variables

SQL Trace Setting

Low		Medium		High	
sql_id	3627211	sql_id	4345330	sql_id	4756333
sql_address	13,106,538,016	sql_address	13,022,460,376	sql_address	13,021,440,688
sql_sid	583	sql_sid	89	sql_sid	312
sql_uid	502	sql_uid	502	sql_uid	502
sql_stmttype	4	sql_stmttype	2	sql_stmttype	4
sql_stmtname	UPDATE	sql_stmtname	SELECT	sql_stmtname	UPDATE
sql_finishtime	1372775249	sql_finishtime	1372775388	sql_finishtime	1372775472
sql_begintxtim	1815720304	sql_begintxtim	1816688410	sql_begintxtim	1817240068
sql_runtime	7.58E-05	sql_runtime	2.66E-06	sql_runtime	9.21E-05
sql_pgreads	0	sql_pgreads	0	sql_pgreads	0
sql_bfreads	4	sql_bfreads	0	sql_bfreads	4
sql_rdcache	100	sql_rdcache	0	sql_rdcache	100
sql_bfidxreads	0	sql_bfidxreads	0	sql_bfidxreads	0
sql_pgwrites	0	sql_pgwrites	0	sql_pgwrites	0
sql_bfwrites	1	sql_bfwrites	0	sql_bfwrites	1
sql_wrcache	100	sql_wrcache	0	sql_wrcache	100
sql_lockreq	5	sql_lockreq	0	sql_lockreq	5
sql_lockwaits	0	sql_lockwaits	0	sql_lockwaits	0
sql_lockwttime	0	sql_lockwttime	0	sql_lockwttime	0
sql_logspace	304	sql_logspace	0	sql_logspace	296
sql_sorttotal	0	sql_sorttotal	0	sql_sorttotal	0
sql_sortdisk	0	sql_sortdisk	0	sql_sortdisk	0
sql_sortmem	0	sql_sortmem	0	sql_sortmem	0
sql_executions	1069	sql_executions	475	sql_executions	504
sql_totaltime	0.118188146	sql_totaltime	0.002672804	sql_totaltime	0.195890748
sql_avgtime	0.00011056	sql_avgtime	5.63E-06	sql_avgtime	0.000388672
sql_maxtime	0.003268819	sql_maxtime	0.001169063	sql_maxtime	0.123698502
sql_numiowait	0	sql_numiowait	0	sql_numiowait	0
sql_avgiowaits	0	sql_avgiowaits	0	sql_avgiowaits	0
sql_totaliowait	0	sql_totaliowait	0	sql_totaliowait	0
sql_rowsperse	13200.98846	sql_rowsperse	376251.0913	sql_rowsperse	10862.90852
sql_estcost	1	sql_estcost	2	sql_estcost	1
sql_estrows	1	sql_estrows	1	sql_estrows	1
sql_actualrows	1	sql_actualrows	0	sql_actualrows	1
sql_sqlerror	0	sql_sqlerror	0	sql_sqlerror	0
sql_isamerror	0	sql_isamerror	0	sql_isamerror	0
sql_isollevel	2	sql_isollevel	2	sql_isollevel	2
sql_sqlmemor	18680	sql_sqlmemor	31064	sql_sqlmemor	18680
sql_numiterat	1	sql_numiterat	1	sql_numiterat	1
sql_database	<None>	sql_database	benchmark3	sql_database	benchmark3
sql_numtables	0	sql_numtables	0	sql_numtables	0
sql_tablelist	None	sql_tablelist	customer	sql_tablelist	district
sql_statement	UPDATE district SET d	sql_statement	SELECT	sql_statement	UPDATE
sql_stmtlen	67	sql_stmtlen	212	sql_stmtlen	67
sql_stmthash	1242825219	sql_stmthash	988199070	sql_stmthash	1242825219
sql_pdq	0	sql_pdq	0	sql_pdq	0
sql_num_hvar	3	sql_num_hvar	3	sql_num_hvar	3
sql_dbspartnu	11534338	sql_dbspartnu	11534338	sql_dbspartnu	11534338
sql_aqt	None	sql_aqt	None	sql_aqt	None
sql_aqtinfo	0	sql_aqtinfo	0	sql_aqtinfo	0

SQL Trace Level differences for Low, Medium, and High

SQL Trace Configuration

- Number of Traces to collect - Default is 1000
- Size of data to collect – Default is 2KB
- Scope of Traces
  - Global – Default is all users
  - User – Specific user list to trace
- **onconfig.std value**

#SQLTRACE
level=low,ntraces=1000,size=2,mode=global

SQL Trace Data

- User ID of the user who ran the command
- Session ID
- Database
- Type of SQL statement
- Duration of the SQL statement execution
- Time statement completed
- Text of the SQL statement or a function call
- Database isolation level

SQL Trace Statistics

- Number of buffer reads and writes
- Number of page reads and writes
- Number of sorts and disk sorts
- Number of lock requests and waits
- Number of logical log records
- Number of index buffer reads
- Estimated number of rows
- Optimizer estimated cost
- Number of rows returned

Sysmaster Tables – View into the SQL Trace Buffer

- Syssqltrace - detailed information about a single SQL statement
- Syssqltrace_info - information about the SQL profile trace system
- Syssqltrace_iter - lists the SQL statement iterators

Syssqltrace -1 of 2

sql_id	Unique SQL execution ID
sql_address	Address of the statement in the code block
sql_sid	Database session ID of the user running the SQL statement
sql_uid	User ID of the statement running the SQL
sql_stmttype	Statement type
sql_stmtname	Statement type displayed as a word
sql_finishtime	Time this statement completed (UNIX)
sql_begintxtime	Time this transaction started
sql_runtime	Statement execution time
sql_pgreads	Number of disk reads for this SQL statement
sql_bfreads	Number of buffer reads for this SQL statement
sql_rdcache	Percentage of time the page was read from the buffer pool
sql_bfidxreads	Number of index page buffer reads
sql_pgwrites	Number of pages written to disk
sql_bfwrites	Number of pages modified and returned to the buffer pool
sql_wrcache	Percentage of time a page was written to the buffer pool
sql_lockreq	Total number of locks required by this SQL statement
sql_lockwaits	Number of times the SQL statement waited on locks
sql_lockwttime	Time the system waited for locks during SQL statement
sql_logspace	Amount of space the SQL statement used in the logical log
sql_sorttotal	Number of sorts that ran for the statement
sql_sortdisk	Number of sorts that ran on disk
sql_sortmem	Number of sorts that ran in memory
sql_executions	Number of times the SQL statement ran
sql_totaltime	Total amount of time spent running the statement
sql_avgtime	Average amount of time spent running the statement
sql_maxtime	Maximum amount of time spent executing the SQL statement

Syssqltrace - 2 of 2

sql_numiowaits	Number of times an I/O operation had to wait
sql_avgiowaits	Average amount of time that the SQL statement had to wait
sql_totaliowaits	Amount of time that the SQL statement had to wait for I/O.
sql_rowspersec	Average number of rows (per second) produced
sql_estcost	Cost associated with the SQL statement
sql_estrows	Estimated number of rows returned for the SQL statement
sql_actualrows	Number of rows returned for the SQL statement
sql_sqlerror	SQL error number
sql_isamerror	RSAM/ISAM error number
sql_isollevel	Isolation level of the SQL statement.
sql_sqlmemory	Number of bytes needed to execute the SQL statement
sql_numiterators	Number of iterators used by the statement
sql_database	Database name
sql_numtables	Number of tables used in executing the SQL statement
sql_tablelist	List of table names directly referenced in the SQL statement.
sql_statement	SQL statement that ran

Display SQL Trace using Onstat –g his

- Shows current setting
- Shows SQL Statements
- Shows Statistics

- Detail displayed depends on Level

SQL API for SQL Trace

- execute function sysadmin:task ("set sql tracing info");
  - The task() function returns a textual message
- execute function sysadmin:admin ("set sql tracing info");
  - The admin() function returns an integer

SQL API for SQL Trace

- "set sql tracing info”
- "set sql tracing off”
- "set sql tracing resume”
- "set sql tracing suspend”
- "set sql tracing on",
  "num_traces","trace_size","level","mode”

SQL API for SQL Trace

- "set sql tracing database add",
  "database_name”
- "set sql tracing database clear”
- "set sql tracing database list”
- "set sql tracing database remove",
  "database_name”

SQL API for SQL Trace

- "set sql tracing session", "clearlofflon", "session_id”
- "set sql tracing user add","user_name”
- "set sql tracing user clear”
- "set sql tracing user list”
- "set sql tracing user remove”
- "set sql user tracing clear","session_id”
- "set sql user tracing off”
- "set sql user tracing on"

SQL Trace Demo

- Examples

Saving SQL Trace Data

- Informix 12.10 – New task ships with the Scheduler called "Save SQL Trace”
- Copies SQL Trace data to the Sysadmin Database
- Need to enable this task in the Sysadmin Database
  - update ph_task set tk_enable = "f" where tk_name = "Save SQL Trace";

Saving SQL Trace Data

- Informix 12.10 – New task ships with the Scheduler called "Save SQL Trace"
- See:
  - $INFORMIXDIR/etc/sysadmin/ sch_sqlcap.sql

SQL Trace Recommendations

- The Sysadmin task to turn on and off SQL trace is more flexible than the ONCONFIG
- Keep the number and size of the SQL Trace buffer small – making the buffer too big will effect Virtual Memory
- Focus on a database or a user
- Save the data for later analysis

Questions?

Send follow-up questions to
lester@advancedatatools.com

Informix Training in 2015

- **Advanced Informix Performance Tuning by Lester Knutsen and Art Kagel**
  - July 20-23, 2015
- **Informix for Database Administrators**
  - October 12-15, 2015–
- 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!

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 Information 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 and Monitoring
Informix Performance Tuning
Informix Training
Informix Consulting
Informix Development

Free Informix Performance Tuning Webcast replays at:

http://advancedatatools.com/Informix/Webcasts.html
YouTube: http://www.youtube.com/AdvancedDataTools

Call: (800) 807-6732 x101 or email: info@advancedatatools.com
or visit us on the web at: http://www.advancedatatools.com

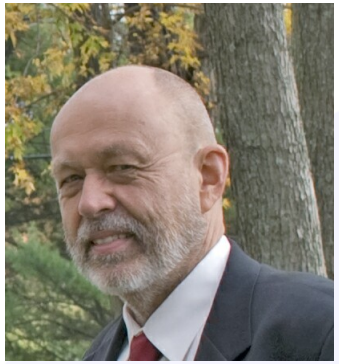

Thank You

Lester Knutsen

Advanced DataTools Corporation

lester@advancedatatools.com

For more information:

http://www.advancedatatools.com