

# **Informix Performance Tuning using the Sysmaster Database by Lester Knutsen**

*Exploring the Sysmaster Database*

***Webcast on April 30, 2013***

# 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@advanceddatatools.com](mailto:lester@advanceddatatools.com)**


**[www.advanceddatatools.com](http://www.advanceddatatools.com)**

**703-256-0267 x102**

***Advanced DataTools***

# Agenda

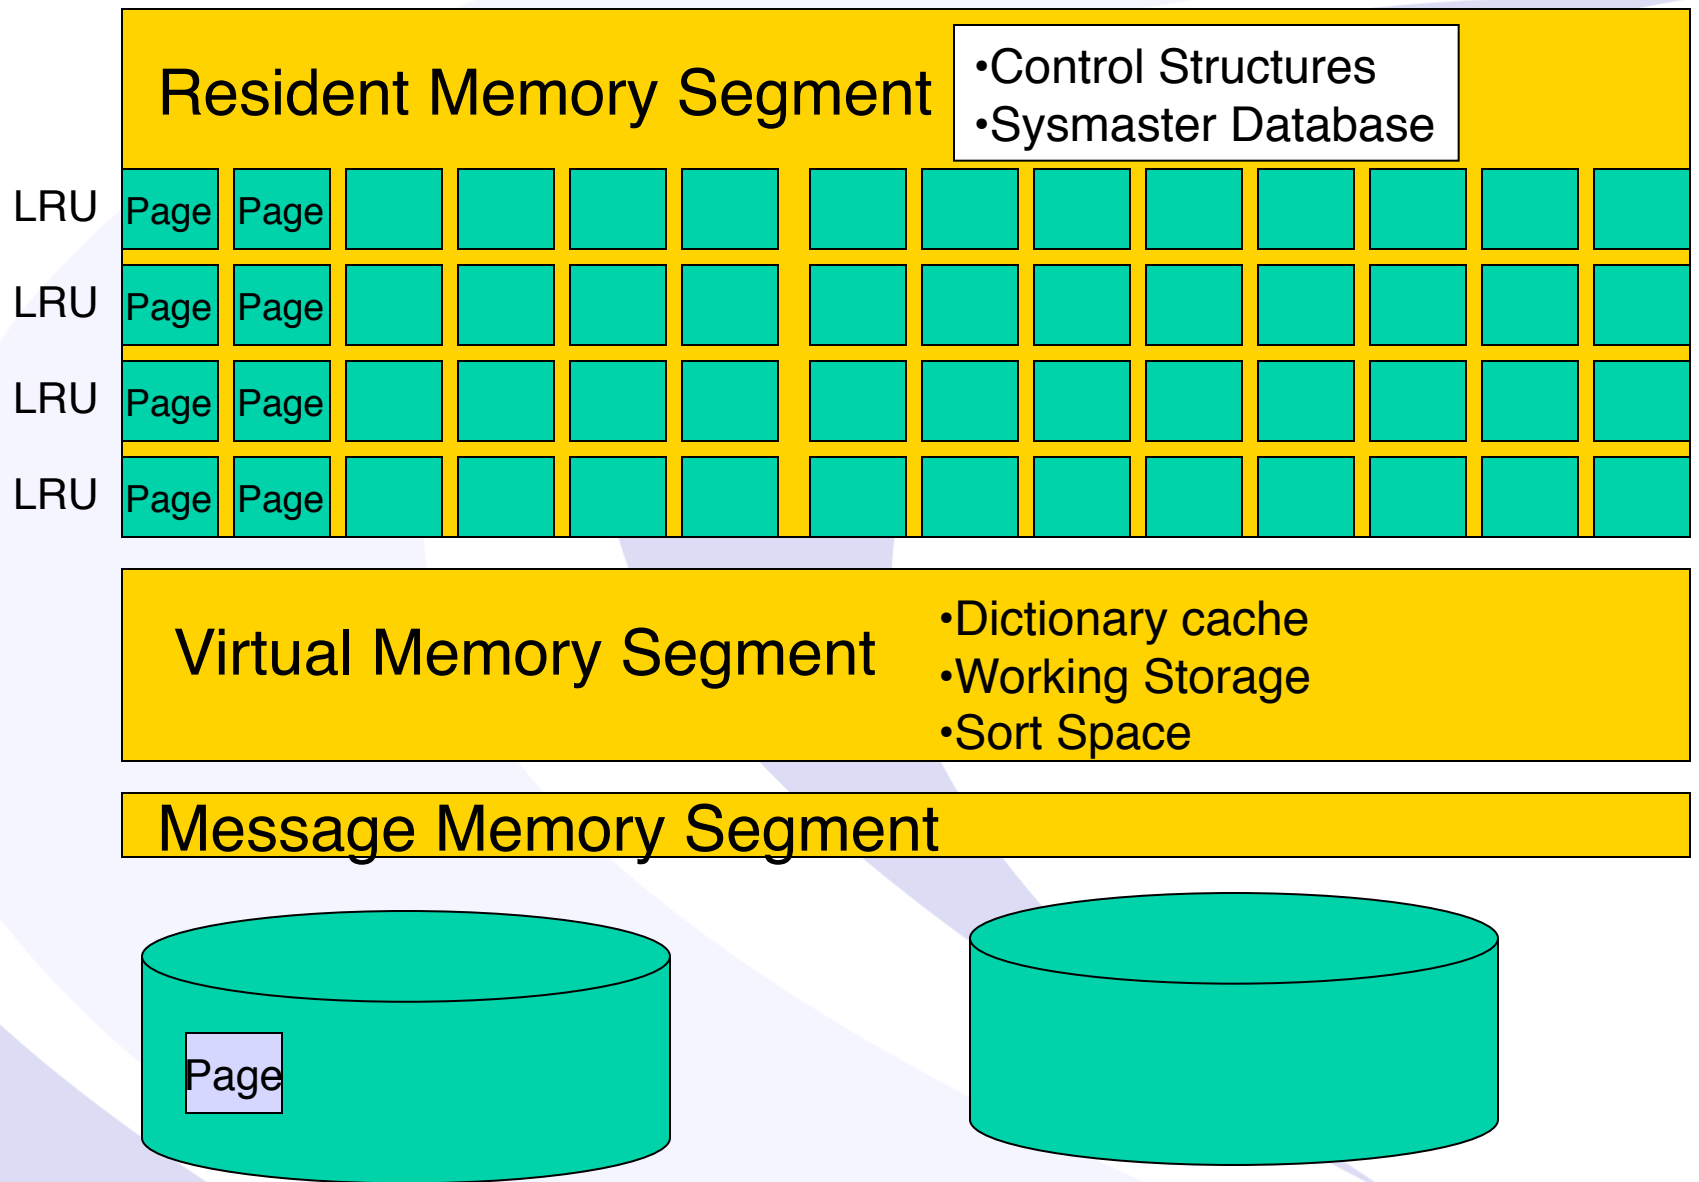
- What is the Sysmaster Database?
- Monitoring your Informix Server
- Monitoring Dbspaces and Chunks
- Monitoring Tables and Indexes
- Monitoring Users and Sessions
- Monitoring SQL



# 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 hid complex data
- Some tables are large (250,000 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

# Using Triggers and Stored Procedures

- Can create triggers and Stored Procedures
- Triggers will never be execute because tables do not change using normal SQL updates
- Use “polling” to check for changes
- Stored Procedures can be executed in the sysmaster database

# 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
- Scripts have been run on versions 7.2 to 12.X
- Sysmaster has changed in IDS 11.X and 12.X

# Interesting table - flags\_text

```
table flags_text
  tabname char(128), -- sysmaster table
  flags   int,      -- flag
  txt     varchar(200) -- description
```

# Server configuration and statistics tables:

- **sysconfig** - ONCONFIG File
- **syslogs** - Logical Logs
- **sysprofile** - Server Statistics
- **sysvpprof** - Virtual Processors

# Sysconfig (onstat -c)

**View sysconfig:** Configuration information from the Informix server. This information was read from the ONCONFIG file when the server was started.

```
cf_id          integer, -- unique numeric identifier
cf_name        char(128), -- config parameter name
cf_flags       integer, -- flags, 0 = in view sysconfig
cf_original    char(513), -- boottime value in ONCONFIG
cf_effective   char(513), -- value effectively in use
cf_default     char(513) -- value by default
```

# What is the current server configuration?

```
select
  cf_name           parameter,
  cf_effective      effective_value
from sysconfig
```

# SQL output

parameter	effective_value
ROOTNAME	rootdbs
ROOTPATH	/u3/dev/rootdbs1
DBSERVERNAME	train1
MIRRORPATH	/u3/dev/rootdbsm1
PHYSDBS	rootdbs
MSGPATH	/u3/informix7/online1.log
CONSOLE	/u3/informix7/console1.log
TAPEDEV	/dev/null
LTAPEDEV	/dev/null
ROOTOFFSET	0
ROOTSIZE	400000



# Syslogs (onstat -l)

## View syslogs: Logical logs status

Number	smallint,	logfile number
Uniqid	integer,	logfile uniqid
Size	integer,	pages in logfile
Used	integer,	pages used in logfile
is_used	integer,	1 for used, 0 for free
is_current	integer,	1 for current
is_backed_up	integer,	1 for backedup
is_new	integer,	1 for new
is_archived	integer,	1 for archived
is_temp	integer,	1 for temp
Flags	smallint	logfile flags

# What is the status of the logical logs?

```
-- List Logical Logs
select
    uniqid,
    used      size_used,
    is_used,
    is_current,
    is_backed_up,
    is_archived
from syslogs
order by uniqid
```

# SQL output

uniqid	size_used	is_used	is_current	is_backed_up	is_archived
32	1000	1	0	1	1
33	1000	1	0	1	1
34	1000	1	0	1	1
35	1000	1	0	1	0
36	1000	1	0	1	0
37	1000	1	0	1	0
38	1000	1	0	1	0
39	1000	1	0	1	0
40	1000	1	0	1	0
41	1000	1	0	1	0
42	1000	1	0	1	0

# Sysprofile (onstat -p)

**View sysprofile:** Current statistics and performance information of the server.

name	char(32),	--profile element name
value	int8	--current value

The values are re-set to 0 when Informix is shutdown and started and when the command “onstat -z” is used.

# Sysprofile – Profile Names

dskreads	ovlock	llgreccs	rapgs_used
bufreads	ovuser	llgpagewrites	seqscans
dskwrites	ovtrans	llgwrites	totalsorts
bufwrites	latchwts	pagreads	memsorts
isamtot	buffwts	pagwrites	disksorts
isopens	lockreqs	flushes	maxsortspace
isstarts	lockwts	compress	ll_niowaits
isreads	ckptwts	fgwrites	ll_iowait_ms
iswrites	deadlks	lruwrites	num_cpu_reac
isrewrites	lktouts	chunkwrites	num_ready
isdeletes	numckpts	btradata	ll_nbfwaits
iscommits	plgpagewrites	btraidx	ll_bfwait_ms
isrollbacks	plgwrites	dpra	

# Sysprofile - onstat -p

Profile

dskreads	pagreads	bufreads	%cached	dskwrits	pagwrits	bufwrits	%cached		
350	372	2108	83.40	10	11	1	0.00		
isamtot	open	start	read	write	rewrite	delete	commit	rollbk	
2839	94	122	457	0	0	0	0	0	
gp_read	gp_write	gp_rewrt	gp_del	gp_alloc	gp_free	gp_curs			
2	0	0	0	0	0	2			
ovlock	ovuserthread	ovbuff	usercpu	syscpu	numckpts	flushes			
0	0	0	3.93	0.87	2	6			
bufwaits	lokwaits	lockreqs	deadlks	dltouts	ckpwaits	compress	seqscans		
67	0	966	0	0	0	0	4		
ixda-RA	idx-RA	da-RA	RA-pgsused	lchwaits					
5	0	247	252	0					

# What are some of the key server statistics?

```
-- Select key Profile values
select name, value from sysprofile
where name in
( "ovlock", "ovuser", "ovtrans",
  "latchwts", "buffwts", "lockwts",
  "ckptwts", "deadlks", "lktouts",
  "fgwrites", "lruwrites",
  "chunkwrites" )
```

# SQL output

name	value
ovlock	0
ovuser	0
ovtrans	0
latchwts	41
buffwts	1617
lockwts	0
ckptwts	12
deadlks	0
lktouts	0
fgwrites	1190
lruwrites	21430
chunkwrites	4648



# What percent of I/O is from buffers?

```
-- Get % read cached
select
  dr.value dskreads, br.value bufreads,
  round ((( 1 - ( dr.value / br.value )) *100 ), 2) cached
from sysprofile dr, sysprofile br
where dr.name = "dskreads"
and   br.name = "bufreads";
-- Get % write cached
select
  dw.value dskwrites, bw.value bufwrites,
  round ((( 1 - ( dw.value / bw.value )) *100 ), 2) cached
from sysprofile dw, sysprofile bw
where dw.name = "dskwrites"
and   bw.name = "bufwrites"
```

# SQL output

dskreads	bufreads	cached
29209	1489235	98.04
dskwrites	bufwrites	cached
56228	414748	86.44

# Sysbuffpool

indx	integer,	index into buffer pool list
address	int8,	address of buffer pool structure
bufsize	integer,	buffer page size in bytes
nbufs	integer,	number of buffers in buffer pool
buff_header	int8,	address of buffer headers
nlrus	smallint,	number of LRUs
mindirty	float,	pct of dirty pages cleaning stops
maxdirty	float,	pct of dirty pages cleaning starts
dskreads	int8,	number of physical block reads
pagreads	int8,	number of pages read from disk
bufreads	int8,	number of buffer cache reads
dskwrites	int8,	number of physical block writes
pagwrites	int8,	number of pages written to disk
bufwrites	int8,	number of buffer cache writes
bufwrites_sinceckpt	int8,	number of buffer cache writes since last checkpoint
bufwaits	int8,	number of buffer waits
ovbuff	int8,	number of buffer pool overflows
flushes	int8,	number of buffer flush calls
fgwrites	int8,	number of foreground writes
lruwrites	int8,	number of lru writes
chunkwrites	int8,	number of chunk writes
lru_time_total	float,	time spent cleaning LRUs
lru_calls	int8	number of times LRU cleaning invoked

# Interesting undocumented table – Sysshmvals\*

sh_mode	int, turbo mode number	sh_optstgbsnum	int, subsystem Blobspace
sh_boottime	int, boot time of day	sh_cpflag	int, TRUE => doing checkpoint
sh_pfcrltime	int, time profilers were last clr	sh_rapages	int, # pages to read ahead
sh_curtime	int, current mt_time	sh_rathreshold	int, # to start next read ahead
sh_bootstamp	int, boot time stamp	sh_lastlogfreed	int, last log (id) written to tape
sh_stamp	int, current time stamp	sh_rmdlktout	int, max timeout when distributed
sh_mainlooptcb	int, address of main thread	sh_narchivers	int, number of active archives
sh_sysflags	int, system operating flags	sh_maxpdqpriority	int, max pdqpriority
sh_maxchunks	int, size of chunk table	sh_fuzcpflag	int, fuzzy checkpoint flag
sh_maxdbspaces	int, size of dbspace table	sh_needcpsyn	int, hard checkpoint
sh_maxuserthreads	int, max # of user structures	sh_nfuzzy	int, # buffers marked fuzzy
sh_maxtrans	int, max # of trans structures	sh_nfuzzypre	int, # buffers fuzzy in last ckpt
sh_maxlocks	int, # of locks total	sh_oldestlsnuq	int, lsn of oldest update not
sh_maxlogs	int, size of log table	sh_oldestlsnpos	int, flushed to disk
sh_nbufs	int, # of buffers total	sh_builddpt	int, builing DPT necessary
sh_pagesize	int, buffer size in bytes	sh_ndptentries	int, # entries in DPT
sh_nlrus	int, # of lru queues	sh_dptsize	int, size of DPT
sh_maxdirty	float, LRU max % dirty pages	sh_curmaxcons	int, max #connections in this run
sh_mindirty	float, LRU min % dirty pages	sh_ovlmaxcons	int, max #connections to server
sh_ncleaners	int, # of cleaning/flushing procs		
sh_longtx	int, # the long transaction flag		

# Using DBINFO with sysmaster

# Time of Server startup:


```
Select DBINFO ('utc_to_datetime', sh_boottime )  
from sysshmvals;
```

# Time Statistics were last cleared (onstat -z) or startup:

```
Select DBINFO ('utc_to_datetime', sh_pfclrtime)  
from sysshmvals;
```

# New - What percent of I/O is from buffers?

```
select bufsize,  
dskreads, pagreads, bufreads,  
round ((( 1-( dskreads / bufreads ) )  
*100 ), 2) read_cach,  
dskwrites, pagwrites, bufwrites,  
round ((( 1-( dskwrites / bufwrites ) )  
*100 ), 2) write_cach  
from sysbufpool;
```



# SQL Output

bufsize	2048
dskreads	989
pagreads	1395
bufreads	18498
read_cach	94.65
dskwrites	191
pagwrites	319
bufwrites	1399
write_cach	86.35

bufsize	4096
dskreads	2
pagreads	4
bufreads	6
read_cach	66.67
dskwrites	1
pagwrites	2
bufwrites	2
write_cach	50.00

# Buffer Turnover Ratio

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



# SQL Output

```
bufsize      2048
pagreads     1395
bufwrites    1399
nbufs        50000
btr          0.0558800000000000
```

```
bufsize      4096
pagreads     4
bufwrites    2
nbufs        40000
btr          0.0001500000000000
```

# Sysvpprof (onstat -g)

**View sysvpprof:** Current statistics on Informix Virtual Processors

```
vpid          integer, -- VP id
txt           char(128)-- VP class name
usecs_user    float, -- unix secs of CPU user time
usecs_sys     float  -- unix secs of CPU system time
```

# What is the status of the virtual processors?

```
-- Select VP Statistics
select
    vpid,
    pid,
    txt[1,5] class,
    round( usecs_user, 2) usercpu,
    round( usecs_sys, 2)  syscpu
from   sysvplst a, flags_text b
where  a.class = b.flags
and    b.tabname = "sysvplst"
```

# SQL output

vpid	pid class	usercpu	syscpu
1	295 cpu	309.26	23.58
2	296 adm	0.14	0.36
3	297 lio	0.27	5.57
4	298 pio	0.15	1.49
5	299 aio	5.00	46.16
6	300 msc	0.04	0.24
7	301 aio	4.65	43.75
8	302 tli	0.14	0.30
9	305 pio	0.22	1.56

# Syscheckpoint table



Intvl	Internal sequence number of the checkpoint
Type	Type of checkpoint, Blocking or Non-Blocking
Caller	Reason for checkpoint, CKPTINT, Physical Log, Logical Log, User
clock_time	Time of checkpoint (System time)
crit_time	Time spent performing checkpoint
flush_time	Time spent flushing pages to disk
cp_time	Time spent from checkpoint pending start to complete
n_dirty_bufs	Number of dirty pages to flush
plogs_per_sec	Average number of pages in physical log per second
llogs_per_sec	Average number of pages in logical log per second
dskflush_per_sec	Average number of disk pages flushed per second
ckpt_logid	Logical Log Id of checkpoint
ckpt_logpos	Logical Log position of checkpoint
Physused	Physical Log pages used
Logused	Logical Log pages used
n_crit_waits	Number of critical waiters threads
tot_crit_wait	Time of critical waiters
longest_crit_wait	Longest Time of crtical waiters
block_time	Time checkpoint blocked threads

# Sysenv and Sysenvses Tables



## **Sysenv**

env_id	Unique numeric identifier
env_name	Environment variable name
env_value	Environment variable value

## **Sysenvses**

envses_sid	Session ID of user
envses_id	Unique numeric identifier
envses_name	Environment variable name
envses_value	Environment variable value

# Sysonlineolog (onstat -m)



<code>offset</code>	Offset into the online log file
<code>next_offset</code>	Offset to the end of the record in the online log file
<code>line</code>	Online log message text

# Sysmgminfo (onstat -g mgm)



max_query	Current maximum number of active PDQ queries allowed
total_mem	Current total memory for PDQ queries
avail_mem	Free PDQ memory
total_seq	Current number of PDQ sequential scans
avail_seq	Unused sequential scans
active	Number of active PDQ queries
ready	Number of PDQ queries waiting and ready to run
min_free_mem	Minimum free PDQ memory
avg_free_mem	Average free PDQ memory
std_free_mem	Standard free PDQ memory
min_free_seq	Minimum free PDQ sequential scans
avg_free_seq	Average free PDQ sequential scans
std_free_seq	Standard free PDQ sequential scans
max_active	Maximum active PDQ sql operations
cnt_active	Count of active PDQ sql operations
avg_active	Maximum active PDQ sql operations
std_active	Standard active PDQ sql operations
max_ready	Maximum ready PDQ sql operations
cnt_ready	Count of ready PDQ sql operations
avg_ready	Average ready PDQ sql operations
std_ready	Standard ready PDQ sql operations



# Results of a Query



Field	Results	Comments
max_query	32	The max number of PDQ queries that can currently run on my laptop is 32
total_mem	512	There is 512 KB of memory available for PDQ queries
avail_mem	0	No memory is available, because it is all used by the one running query
total_seq	1048576	The max number of sequence scans that can be run
avail_seq	1048575	The number that is currently available
active	1	Number of active PDQ queries currently running
ready	0	Number of PDQ queries waiting to run. This is a key field and will inform you if you have a backlog of queries ready but waiting on a resource to run.
min_free_mem	0	Minimum free memory
avg_free_mem	0	Average free memory
std_free_mem	0	Standard free memory
min_free_seq	1048575	Minimum free scans
avg_free_seq	1048575	Average free scans
std_free_seq	0	Standard free scans
max_active	1	Maximum number of active queries
cnt_active	4	Count of number of queries that have been run since the server was started
avg_active	1	Average number of active queries
std_active	0	Standard number of active PDQ queries
max_ready	0	Maximum number of queries waiting and ready to run
cnt_ready	0	Count of number of queries that have had to wait since the server was started
avg_ready	0	Average number of waiting and ready queries
std_ready	0	Standard number of waiting and ready queries



# Sysnetclienttype, Sysnetglobal, Sysnetworkio

- Sysnetclienttype - an overview of the network activity for each client type
- Sysnetglobal - provides an overview of the system network
- Sysnetworkio - provides an overview of the system network I/O

# Sysnetclienttype



nc\_cons\_allowed  
nc\_accepted  
nc\_rejected  
nc\_reads  
nc\_writes  
nc\_name

Are connections allowed  
Number of connections  
Number of cons rejected  
Number of network reads  
Number of network writes  
Name

# Sysnetglobal



<code>ng_reads</code>	Number of network reads
<code>ng_writes</code>	Number of network writes
<code>ng_his_read_count</code>	Network reads by discon users
<code>ng_his_write_count</code>	Network writes by discon users
<code>ng_his_reads_bytes</code>	Amount of data txfr by discon users
<code>ng_his_writes_bytes</code>	Amount of data txfr by discon users
<code>ng_connects</code>	Number of connections
<code>ng_num_netscbs</code>	Number of netscbs
<code>ng_max_netscbs</code>	Max number of netscbs
<code>ng_free_thres</code>	Free threshold
<code>ng_free_cnt</code>	Free count
<code>ng_wait_thres</code>	Wait threshold
<code>ng_wait_cnt</code>	Wait count
<code>ng_pvt_thres</code>	
<code>ng_netbuf_size</code>	Network buffer size
<code>ng_buf_alloc</code>	
<code>ng_buf_alloc_max</code>	
<code>ng_netscb_id</code>	Next netscb id

# Sysnetworkio



net_id	Network ID
sid	Session id
net_netscb	Address of netscb
net_client_type	Client type
net_client_name	Client protocol name
net_read_cnt	Number of read operations
net_read_bytes	Number of bytes txfr to server
net_write_cnt	Number of write operations
net_write_bytes	Number of bytes txfr to client
net_open_time	Time connection was made
net_last_read	Time of last network read
net_last_write	Time of last network write
net_state	State of network connection
net_options	Sqlhost options
net_prot_id	ID for protocol name
net_protocol	Protocol name
net_server_fd	Poll server fd
net_poll_thread	Poll thread id

## Dbospace & chunk tables:

- sysdbspaces - DB Spaces
- syschunks - Chunks
- syschkio - I/O by Chunk
- syschfree\* - Free Space by Chunk

# Sysdbspaces (onstat -d)

**View sysdbspaces:** List all dbspaces on the server

```
Dbsnum      smallint, -- dspace number,
name        char(128), -- dspace name,
owner       char(32), -- dspace owner,
pagesize  int, -- page size in Informix 10.X
Fchunk     smallint, -- first chunk in dspace,
nchunks    smallint, -- number of chunks in dspace,
is_mirrored bitval, -- dspace mirrored, 1=Yes, 0=No
is_blobspace bitval, -- dspace a blob space, 1=Yes
is_temp     bitval, -- dspace temp, 1=Yes, 0=No
flags      smallint -- dspace flags
```

# Syschunks (onstat -d)

**View syschunks:** Lists all chunks on the server

```
Chknum      smallint, -- chunk number
Dbsnum      smallint, -- dbspace number
nxchknum    smallint, -- number of next chunk
pagesize  smallint, -- page size in Informix 10.X
chksize     integer, -- pages in chunk
offset      int8, -- pages offset into device
nfree       integer, -- free pages in chunk
is_offline  bitval, -- chunk offline, 1=Yes, 0=No
is_recovering bitval, -- chunk recovering, 1=Yes
is_blobchunk bitval, -- chunk blobchunk, 1=Yes
is_inconsistent bitval, -- chunk inconsistent, 1=Yes
```



# Syschunks (continued)

```
flags          smallint, -- flags converted by bitval
fname          char(256), -- device pathname
mfname        char(256), -- mirror device pathname
moffset        integer, -- pages offset into mirror
mis_offline    bitval, -- mirror chunk offline, 1=Yes
mis_recovering bitval, -- mirror chunk recovering,
mflags         smallint  mirror chunk flags
```

# Syschkio (onstat -D)

**View syschkio:** Lists I/O statistics by chunk

```
chunknum      smallint, -- chunk number
reads         integer, -- number of read ops
pagesread    integer, -- number of pages read
writes        integer, -- number of write ops
Pageswritten  integer, -- number of pages written
mreads        integer, -- number of mirror read ops
mpagesread   integer, -- number of mirror pages read
mwrites       integer, -- number of mirror write ops
mpageswritten integer -- number of mirror pages written
```

**Note - Underlying undocumented table has read and write time by chunk - see table syschktab\_fast**

# Syschfree\*

**Table syschfree:** Lists free space on a chunk

chknum	integer, -- chunk number
extnum	integer, -- extent number in chunk
start	integer, -- physical addr of start
leng	integer -- length of extent

# How much dbspace is free?

```
-- dbsfree.sql
Select d.dbsnum,
       name dbspace,
       sum(chksize) Pages_size, -- sum of all chunk pages
       sum(chksize) - sum(nfree) Pages_used,
       sum(nfree) Pages_free, -- sum of all chunks free pages
       round ((sum(nfree)) / (sum(chksize)) * 100, 2)
       Percent_free
from sysdbspaces d, syschunks c
where d.dbsnum = c.dbsnum
and d.is_blobspace = 0
group by 1, 2
order by 1;
```

# SQL output

dbspace	pages_size	pages_used	pages_free	percent_free
rootdbs	20000	5653	14347	71.74
logsdbs	12500	12053	447	3.58
datadbs	25000	6722	18278	73.11
tmpdbs	12500	53	12447	99.58

# How much blob space is free?

```
-- blobfree.sql
select
  name dbspace,
  sum(chksize) Size_in_Pages, -- sum of all chunks
  sum(nfree) Num_free_blob_page -- sum of freepages
from   sysdbspaces d, syschunks c
where  d.dbsnum = c.dbsnum
and    d.is_blobspace = 1
group by 1
order by 1
```

# SQL output

dbspace	size_in_pages	num_free_blob_page
blobdbs	10000	2497

# Where are blocks of free dbspace?

```
-- chkflist.sql
select
  name dbspace, -- dbspace name
  f.chknum, -- chunk number
  f.extnum, -- extent number of free space
  f.start, -- starting address of free space
  f.leng free_pages -- length of free space
from sysdbspaces d, syschunks c, syschfree f
where d.dbsnum = c.dbsnum
and c.chknum = f.chknum
order by dbspace, free_pages desc
```



# SQL output

dbspace	chknun	extnum	start	free_pages
datadbs	4	0	3	12497
datadbs	3	31	9107	3393
datadbs	3	15	1921	976
datadbs	3	13	1705	160
datadbs	3	30	6069	160
datadbs	3	25	5429	128
datadbs	3	19	4853	96
datadbs	3	29	5909	96
datadbs	3	24	5333	64
datadbs	3	26	5621	64

# What chunks have the most I/O?

```
-- chkio.sql
select name
       dbspace,
       chknum, "Primary" chktype,
       reads, writes,
       pagesread, pageswritten
from   syschktab c, sysdbstab d
where  c.dbsnum = d.dbsnum
union all
select
       name  dbspace,
       chknum, "Mirror"      chktype,
       reads, writes,
       pagesread, pageswritten
from   sysmchktab c, sysdbstab d
where  c.dbsnum = d.dbsnum
order by 1,2,3;
```

# SQL output

dbspace	chknun	chktype	reads	writes	pagesread	pageswritten
blobdbs	6	Primary	21	3	31	10
datadbs	3	Primary	2082	31	9087	31
datadbs	4	Primary	5	0	7	0
logsdbs	2	Primary	176	996	1347	11704
rootdbs	1	Mirror	11616	26196	22499	30102
rootdbs	1	Primary	13340	26111	22271	30102
tmpdbs	5	Primary	13	2	13	3

# What is the status of chunks?

```
select      name dbspace,          -- dbspace name
            d.dbsnum,             -- dbspace num
            is_mirrored,          -- dbspace is mirrored 1=Yes 0=No
            is_blobspace,        -- dbspace is blobspace 1=Yes 0=No
            is_temp,              -- dbspace is temp 1=Yes 0=No
            chunknum chunknum,    -- chunk number
            fname device,         -- dev path
            offset dev_offset,    -- dev offset
            is_offline,           -- Offline 1=Yes 0=No
            is_recovering,        -- Recovering 1=Yes 0=No
            is_blobchunk,         -- Blobspace 1=Yes 0=No
            is_inconsistent,      -- Inconsistent 1=Yes 0=No
            chksize Pages_size,   -- chunk size in pages
            nfree Pages_free,     -- chunk free pages
            mfname mirror_device, -- mirror dev path
            mis_recovering_offse  -- mirror recovering 1=Yes 0=No
from        sysdbspaces d, syschunks c
where       d.dbsnum = c.dbsnum
order by   dbsnum, dbspace, chunknum
```

# SQL output

dbspace	rootdbs
dbnum	1
is_mirrored	1
is_blobspace	0
is_temp	0
chunknum	1
device	/u3/dev/rootdbs1
dev_offset	0
is_offline	0
is_recovering	0
is_blobchunk	0
is_inconsistent	0
pages_size	20000
pages_free	14355
mirror_device	/u3/dev/rootdbsm1
mirror_offset	0

# Database & table information tables:

- sysdatabases - Databases
- systabnames - Tables
- sysextents - Tables extents
- sysptprof - Tables I/O
- systabinfo\*  
information - Tables

# Sysdatabases

**View sysdatabases:** List of databases on the server.

```
name          char(128), -- database name
partnum       integer,  -- table id for systables
owner         char(32), -- user name of creator
created       integer, -- date created
is_logging    bitval,  -- unbuffered logging, 1=Yes, 0=No
is_buff_log   bitval,  -- buffered logging, 1=Yes, 0=No
is_ansi       bitval,  -- ANSI mode database, 1=Yes, 0=No
is_nls        bitval,  -- NLS support, 1=Yes, 0=No
flags         smallint -- logging flags
```

# Systabnames

**Table systabnames:** All tables on the server.

```
partnum      integer, -- table id for table
dbsname      char(128), -- database name
owner        char(32), -- table owner
tablename    char(128), -- table name
collate      char(32) -- collation associated
                with NLS DB
```



# Sysexts (oncheck -pe)

**View sysexts:** Tables and each extent on the server.

dbsname	char(128), -- database name
tabname	char(128), -- table name
Start	integer, -- extent physical address
size	integer -- size of this extent

# Sysptprof

## View sysptprof: Tables IO profile.

```
dbname          char(128), -- database name
tablename       char(128), -- table name
partnum         integer, -- partnum for this table
lockreqs        integer, -- lock requests
lockwts         integer, -- lock waits
deadlks         integer, -- deadlocks
lktouts         integer, -- lock timeouts
isreads         integer, -- reads
iswrites        integer, -- writes
isrewrites      integer, -- rewrites
isdeletes       integer, -- deletes
bufreads        integer, -- buffer reads
bufwrites       integer, -- buffer writes
seqscans        integer, -- sequential scans
pagreads        integer, -- disk reads
pagwrites       integer-- disk writes
```

# Systabinfo\*

## View systabinfo: Table information

ti_partnum	integer, -- table's partnum
ti_flags	smallint, -- partition flags
ti_rowsize	smallint, -- rowsize (max for variable)
ti_ncols	smallint, -- number of varchar or blob columns
ti_nkeys	smallint, -- number of indexes
ti_nextns	smallint, -- number of extents
ti_created	integer, -- date created
ti_serialv	integer, -- current serial value
ti_fextsiz	integer, -- first extent size ( in pages )
ti_nextsiz	integer, -- next extent size ( in pages )
ti_nptotal	integer, -- number of pages allocated
ti_npused	integer, -- number of pages used
ti_npdata	integer, -- number of data pages
ti_octptnm	integer, -- OCT partnum (optical blobs only)
ti_nrows	integer -- number of data rows

# What databases are on the server?

```
-- dblist.sql
select
-- use dbinfo function to convert partnum to
  dbspace
  dbinfo("DBSPACE",partnum) dbspace,
  name database,
  owner,
  is_logging,
  is_buff_log
from sysdatabases
order by dbspace, name;
```

# SQL output

dbspace	database	owner	is_logging	is_buff_log
datadbs	extentdb2	usr2	0	0
datadbs	zip1	usr1	0	0
datadbs	zip_lk	lester	0	0
rootdbs	extentdb	lester	0	0
rootdbs	extentdb1	usr1	0	0
rootdbs	onpload	lester	1	0
rootdbs	stores1	usr1	0	0
rootdbs	stores2	usr2	0	0
rootdbs	stores7	informix	0	0
rootdbs	sysmaster	informix	1	0

# What is the size of my databases?

```
select  dbsname,  
        sum( pe_size ) total_pages  
from    systabnames, sysptnext  
where   partnum = pe_partnum  
group  by 1  
order  by 2 desc
```

# What tables have extents?

```
-- tabextents.sql
select
    dbsname,
    tabname,
    count(*)          num_of_extents,
    sum( pe_size )   total_size
from  systabnames, sysptnext
where partnum = pe_partnum
group by 1, 2
order by 3 desc, 4 desc;
```

# SQL output

dbsname	tablename	num_of_extents	total_size
zip7	zip	50	1168
zip_lk	zip	27	1544
rootdbs	TBLSpace	8	400
sysmaster	syscolumns	6	56
datadbs	TBLSpace	4	200
sysmaster	sysviews	3	24
sysmaster	sysprocbody	3	24
sysmaster	systables	3	24
extentdb1	extent_sizes	2	24
sysutils	sysprocbody	2	16
sysmaster	sysconstraints	2	16
stores2	sysprocbody	2	16



# How calculate new extent sizes?

```
-- tabextprop.sql
select  dbsname,
        tabname,
        count(*) num_of_extents,
        sum (pe_size ) current_pages_used,
        round (sum (pe_size )
                * 2 { Your systems page size in KB }
                * 1.2 { Add 20% Growth factor })
        Proposed_ext_size, { First Extent Size in KB }
        round (sum (pe_size )
                * 2 { Your systems page size in KB }
                * .2 { Estimated 20% Yearly Growth })
        Proposed_next_size { Next Extent Size in KB }
from    systabnames, sysptnext
where   partnum = pe_partnum
group  by 1, 2
order  by 3 desc, 4 desc;
```

# SQL output

dbname	zip7
tablename	zip
num_of_extents	50
current_pages_used	1168
proposed_ext_size	2803
proposed_next_size	467

dbname	zip_lk
tablename	zip
num_of_extents	27
current_pages_used	1544
proposed_ext_size	3706
proposed_next_size	618

# What tables have the most I/O?

```
-- tabprofile.sql
select
  dbname,
  tablename,
  DBINFO ( 'dbspace', partnum ),
  lockreqs, lockwts, deadlks, lktouts,
  isreads, iswrites, isrewrites, isdeletes,
  bufreads, bufwrites, seqscans, pagreads,
  pagwrites
from sysptprof
order by isreads desc;
-- change this sort to whatever you need to
monitor.
```

# SQL output

dbsname	tablename	isreads	iswrites	lockreqs
zip	zip	41898	41898	830
sysmaster	systables	11402	0	67187
sysmaster	sysusers	10276	315	51373
sysmaster	sysviews	2653	0	15919
sysmaster	sysprocauth	2212	0	13272
zip_lk	zip	1399	0	1
sysmaster	sysprocedures	1108	0	6649
sysmaster	syscolumns	872	0	5182
sysmaster	sysdatabases	538	3	1469
sysmaster	flags_text	450	0	2546
zip	systables	101	34	239
sysmaster	systabauth	86	0	536

# What tables have sequence scans?

```
select  dbsname,  
        tabname,  
        sum(seqscans) total_scans  
from    sysptprof  
where   seqscans > 0  
group  by 1, 2  
order  by 3 desc
```

# User session information tables:

- **sysessions** - Session data
- **sysesprof** - User statistics
- **syslocks** - Locks
- **syseswts** - Wait times

# Sysessions (onstat -g ses)

**View sysessions:** User session and connection information

sid	integer,	-- Session id number
username	char(32),	-- User name
uid	smallint,	-- User unix id
pid	integer,	-- User process id
hostname	char(16),	-- Hostname
tty	char(16),	-- TTY port
connected	integer,	-- Time user connected
feprogram	char(16),	-- Program name
pooladdr	integer,	-- Pointer to private session pool

# Sysessions (continued)

```
is_wlatch    integer, -- Flag 1=Yes wait on latch
is_wlock     integer, -- Flag 1=Yes wait on lock
is_wbuff     integer, -- Flag 1=Yes wait on buffer
is_wckpt     integer, -- Flag 1=Yes wait on checkpoint
is_wlogbuf   integer, -- Flag 1=Yes wait on log buffer
is_wtrans    integer, -- Flag 1=Yes wait on a transaction
is_monitor   integer, -- Flag 1=Yes a monitoring process
is_incrit    integer, -- Flag 1=Yes in crtical section
State        integer  -- Flags
```



# Sysprof (onstat -g ses)

## View sysprof: User session performance statistics

sid	integer, -- Session Id
lockreqs	decimal(16,0), -- Locks requested
locksheld	decimal(16,0), -- Locks held
Lockwts	decimal(16,0), -- Locks waits
Deadlks	decimal(16,0) -- Deadlocks detected
Lktouts	decimal(16,0), -- Deadlock timeouts
Logrecs	decimal(16,0), -- Logical Log written
Isreads	decimal(16,0), -- Reads
iswrites	decimal(16,0), -- Writes
isrewrites	decimal(16,0), -- Rewrites
isdeletes	decimal(16,0), -- Deletes
iscommits	decimal(16,0), -- Commits
isrollbacks	decimal(16,0), -- Rollbacks
longtxs	decimal(16,0), -- Long transactions

# Sysseprof (continued)

bufreads	decimal(16,0),	-- Buffer reads
bufwrites	decimal(16,0),	-- Buffer writes
seqscans	decimal(16,0),	-- Sequential scans
pagreads	decimal(16,0),	-- Page reads
pagwrites	decimal(16,0),	-- Page writes
total_sorts	decimal(16,0),	-- Total sorts
dsksorts	decimal(16,0),	-- Sorts to disk
max_sortdiskspace	decimal(16,0),	-- Max space used by a sort
logspused	decimal(16,0),	-- Current log bytes used
maxlogsp	decimal(16,0)	-- Max bytes of logical logs used

# Syslocks (onstat -k)

## **View syslocks:** Active locks on server

dbsname	char(128),	-- Database name
tabname	char(128),	-- Table name
rowidlk	integer,	-- Rowid for index key lock
Keynum	smallint,	-- Key number of index key lock
owner	integer,	-- Session ID of lock owner
Waiter	integer,	-- Session ID of first waiter
type	char(4),	-- Type of Lock

# Syslocks – Type of Locks

B - byte lock

IS - intent shared lock

S - shared lock

XS - repeatable read shared key

U - update lock

IX - intent exclusive lock

SIX - shared intent exclusive

X - exclusive lock

XR - repeatable read exclusive

# Syseswts

## **View syssteswts:** Wait status and times on objects

```
sid      integer, -- Session ID
reason  char(50), -- Description of reason for wait
numwaits integer, -- Number of waits for this reason
Cumtime float, -- Cumulative wait time for this reason
Maxtime integer -- Max wait time for this reason
```

# Example SQL: dbwho.sql

```
select sysdatabases.name database, -- Database Name
       syssessions.username,      -- User Name
       syssessions.hostname,      -- Workstation
       syslocks.owner sid         -- Informix Session ID
from   syslocks, sysdatabases , outer syssessions
Where  syslocks.tabname = "sysdatabases"
       -- Locks on sysdatabases
and    syslocks.rowidlk = sysdatabases.rowid
       -- Join to database
and    syslocks.owner = syssessions.sid
       -- Use session ID
order by 1;
```

# Dbwho shell script

```
#!/bin/sh
# Program: dbwho
# Description: List database, user and workstation
echo "Generating list of users by database ..."
dbaccess sysmaster - <<EOF
Select sysdatabases.name database,
        syssessions.username,
        syssessions.hostname,
        syslocks.owner sid
from   syslocks, sysdatabases , outer syssessions
where  syslocks.rowidlk = sysdatabases.rowid
and    syslocks.tabname = "sysdatabases"
and    syslocks.owner = syssessions.sid;
order by 1;
EOF
```

# List all Active Sessions

```
-- sessions.sql
select  sid,
        username,
        pid,
        hostname,
        l2date(connected) startdate
        -- convert unix time to date
from    syssessions
```

## Sample Output

sid	username	pid	hostname	startdate
47	lester	11564	merlin	07/14/1997



# List Users Waiting on Resources

```
-- seswait.sql
select username,
       is_wlatch,  -- blocked waiting on a latch
       is_wlock,  -- blocked waiting on a locked record or table
       is_wbuff,  -- blocked waiting on a buffer
       is_wckpt,  -- blocked waiting on a checkpoint
       is_incrit  -- session is in a critical section of
                  -- transaction (e.g writting to disk)

from   syssessions
order by username;
```

## Sample Output

username	is_wlatch	is_wlock	is_wbuff	is_wckpt	is_incrit
lester	0	1	0	0	0
lester	0	0	0	0	0
lester	0	0	0	0	0

# Monitor Resource Usage by User

```
-- sesprof.sql
select      username,
            sysprof.sid,
            lockreqs,
            bufreads,
            bufwrites
from        sysprof, syssessions
where      sysprof.sid = syssessions.sid
order by   bufreads desc
```

# Syssqltrace, Syssqltrace\_info, Syssqltrace\_iter



- 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_lockwtttime	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_numioawaits	Number of times an I/O operation had to wait
sql_avgioawaits	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

# Systhreads – View of systcblst

th_id	thread id
th_addr	thread address
th_joinlist	head of joined threads list
th_joinnext	next thread of joined threads list
th_joinee	thread this thread joined
th_name	thread name
th_state	thread state
th_priority	thread priority
th_class	thread class
th_vpid	VP where thread is running
th_mtxwait	pointer to mutex thread is waiting on
th_conwait	pointer to condition thread is waiting on
th_waketime	seconds slept + start time
th_startwait	sleep start time
th_startrun	total time thread has run



# SQL Explain Output

## View Syssexplain:

sqx\_sessionid, sqx\_sdbno, sqx\_iscurrent,  
sqx\_executions, sqx\_cumtime, sqx\_bufreads,  
sqx\_pagereads, sqx\_bufwrites, sqx\_pagewrites,  
sqx\_totsorts, sqx\_dsksorts, sqx\_sortspmax,  
sqx\_conbno, sqx\_ismain, sqx\_selflag,  
**sqx\_estcost**, sqx\_estrows, sqx\_seqscan,  
sqx\_srtscan, sqx\_autoindex, sqx\_index,  
sqx\_remsql, sqx\_mrgjoin, sqx\_dynhashjoin,  
sqx\_keyonly, sqx\_tempfile, sqx\_tempview,  
sqx\_secthread, sqx\_sqlstatement

# How to find expensive queries running now?

```
Select sqx_estcost,  
       sqx_sessionid,  
       sqx_sqlstatement  
from syssexplain  
order by sqx_estcost desc
```

This will list all current running queries ordered by the SQL Explain estimated costs



# Next Webcast

- Date: June 4<sup>th</sup>, 2013
- Time: 2:00pm EST
- Informix Remote DBA Monitoring
  - What you need to monitor
  - Building a Data Warehouse for monitoring
  - Building a Dashboard for monitoring using Yellowfin BI

# Fastest Informix DBA Contest 2013

- This year the Fastest Informix DBA Contest will be held over the web so anyone, anywhere in the world can participate.
- The challenge will be a combination of OLTP and batch. Who can get the most transactions per minute in an OLTP benchmark and at the same time run a batch billing job and generate the most bills.
- Contest details and challenge will be posted June 1, 2013. Watch here for more details.

<http://www.advanceddatatools.com/Informix/index.html>

# Informix Training in 2013

- July 15-18, 2013
  - **Advanced Informix Performance Tuning**
- September 9-12, 2013
  - **Informix for Database Administrators**
- October 28-31, 2013 –
  - **Advanced Informix Performance Tuning**
- 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@advanceddatatools.com](mailto:lester@advanceddatatools.com)



Thank You

Lester Knutsen  
***Advanced DataTools Corporation***

lester@advanceddatatools.com

For more information:

<http://www.advanceddatatools.com>

***Advanced DataTools***