
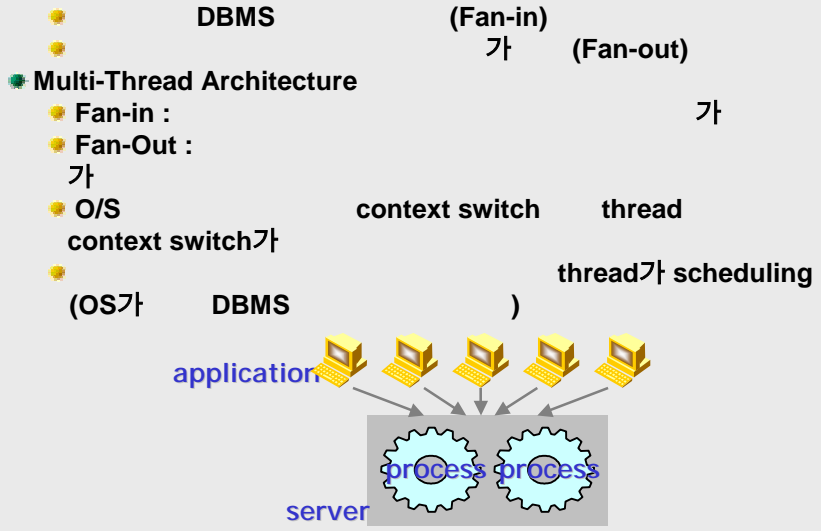


1. *INFORMIX ARCHITECTURE*

The Informix Database Server
System Architecture Overview
Process Component
Shared Memory Component
Disk Component

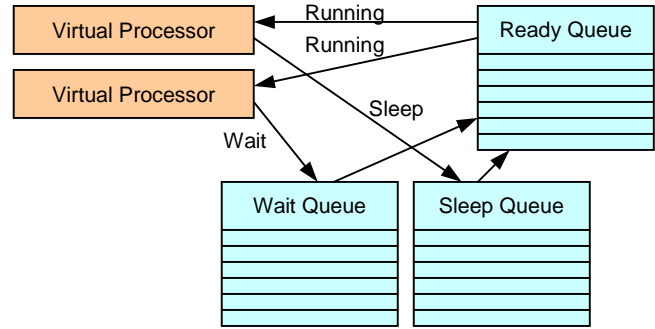
The Informix Database Server

Multi-Thread Architecture DBMS



User request

(dynamic) 가 Dynamic Server 가 (scalability)

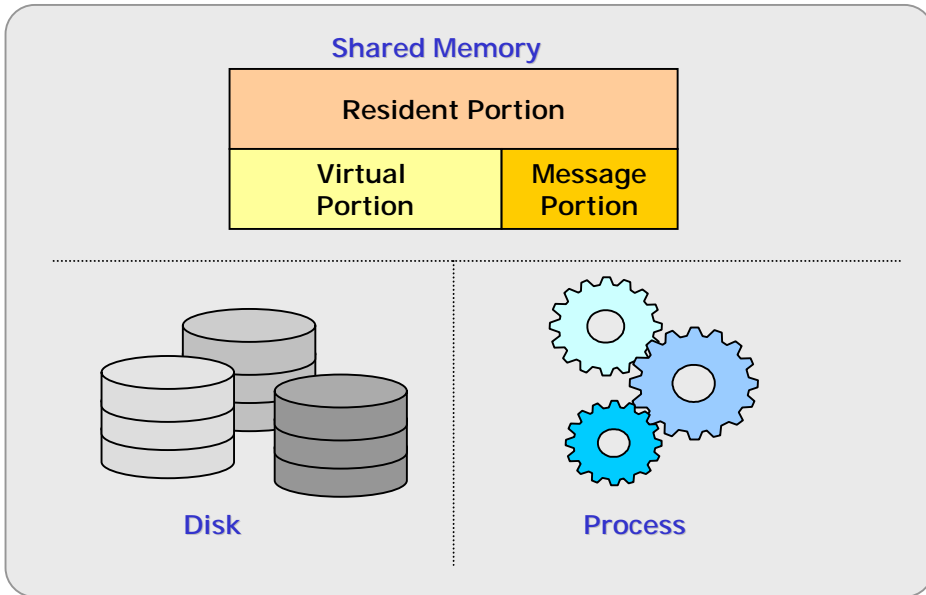


Informix> onstat -g ath

Threads:

tid	tcb	rstcb	prty	status	vp-class	name
2	a76ce68	0	2	sleeping forever	3lio	lio vp 0
3	a806148	0	2	sleeping forever	4pio	pio vp 0
4	a81b148	0	2	sleeping forever	5aio	aio vp 0
5	a830148	0	2	sleeping forever	6msc	msc vp 0
6	a85d148	0	2	sleeping forever	7aio	aio vp 1
7	a872238	a30c018	4	sleeping secs: 1	1cpu	main_loop()
8	a872d18	0	2	running	1cpu	tlitcpgoll
9	a806800	0	2	running	8shm	sm_poll
10	a8a7768	0	3	sleeping forever	1cpu	tlitcplst
11	a830308	0	3	sleeping forever	1cpu	sm_listen
12	a85dbf0	0	2	sleeping secs: 1	1cpu	sm_discon
13	a85ddb0	a30c608	2	sleeping forever	1cpu	flush_sub(0)
14	a8f5018	a30cbf8	2	sleeping forever	1cpu	flush_sub(1)
15	a8f52b0	0	2	sleeping forever	9aio	aio vp 2
16	a8f5470	0	2	sleeping forever	10aio	aio vp 3
17	a92a148	0	2	sleeping forever	11aio	aio vp 4
18	a93f148	0	2	sleeping forever	12aio	aio vp 5
19	a954378	a30d1e8	3	sleeping forever	1cpu	aslogflush
20	a9548f8	a30d7d8	2	sleeping secs: 42	1cpu	btclean
37	a8f5bd8	a30e3b8	4	sleeping secs: 1	1cpu	onmode_mon
43	a8f5e70	a30ddc8	2	cond wait netnrm	1cpu	sqllexc

System Architecture Overview



```
Informix> onstat -g seg

Segment Summary:
id      key          addr      size      ovhd      class blkused blkfree
2162    1388857345  a000000  2097152  215344    R      274     238
2163    1388857346  a200000  8388608  856       V      1978    70
373     1388857347  aa00000  1048576  632       M      133     123
374     1388857348  ab00000  8388608  856       V      1       2047
Total:  -          -          19922944 -          -      2386    2478
```

(* segment locked in memory)

• Process Component


- CPU가 scheduling, running thread scheduling, running 가 (virtual processor)
- VP virtual processor class
- VP root OS resource user가 (kill -9) . root informix id

• Shared Memory Component

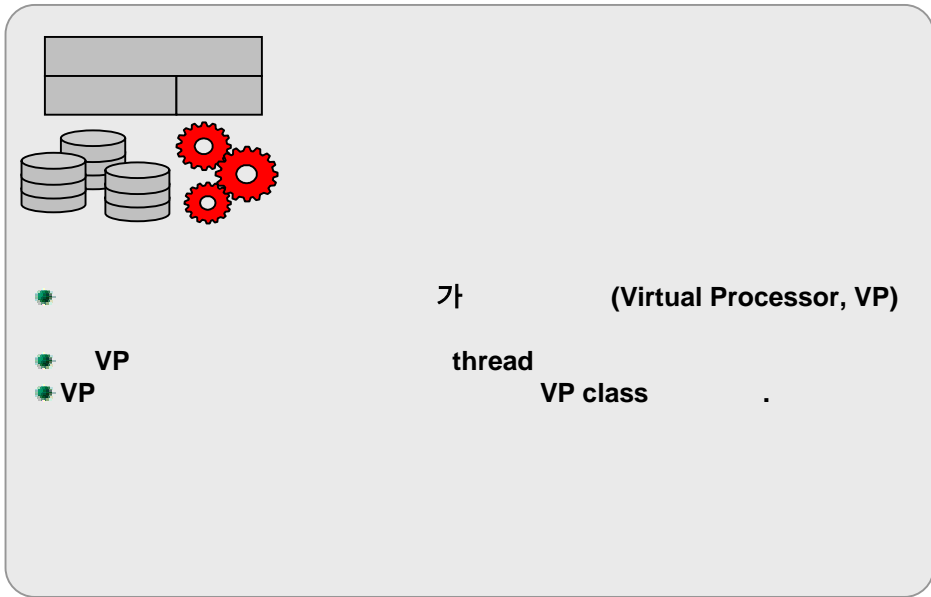
- 가
- resident portion
 - 가
 - (lock, log, buffer) internal
 - table 가
 - internal table
 - onmonitor onconfig
- virtual portion
 - User thread
 - /
- message portion
 - client server가

• Disk Component

- Physical space logical space
- Physical space : 가
 - chunk, extent, page
- Logical space : 가
 - dbspace(blobspace, sbspace), tblspace



Process Component



Virtual Processor Class

CPU	thread 가	
PIO	physical log thread	1 or 2 (mirror)
LIO	Logical log thread	1 or 2 (mirror)
AIO	Cooked chunk I/O kernel	
ADT	auditing thread	1 or 2 (mirror)
MSC	miscellaneous task	1
SHM	thread	
TLI	TLI thread	
SOC	Socket thread	
ADM	timer	1
OPT	optical BLOB	0 or 1
JVP	JAVA UDR JVM(Java virtual machine)	
Class	(UDR) VP	

```
Informix> onstat -g glo
```

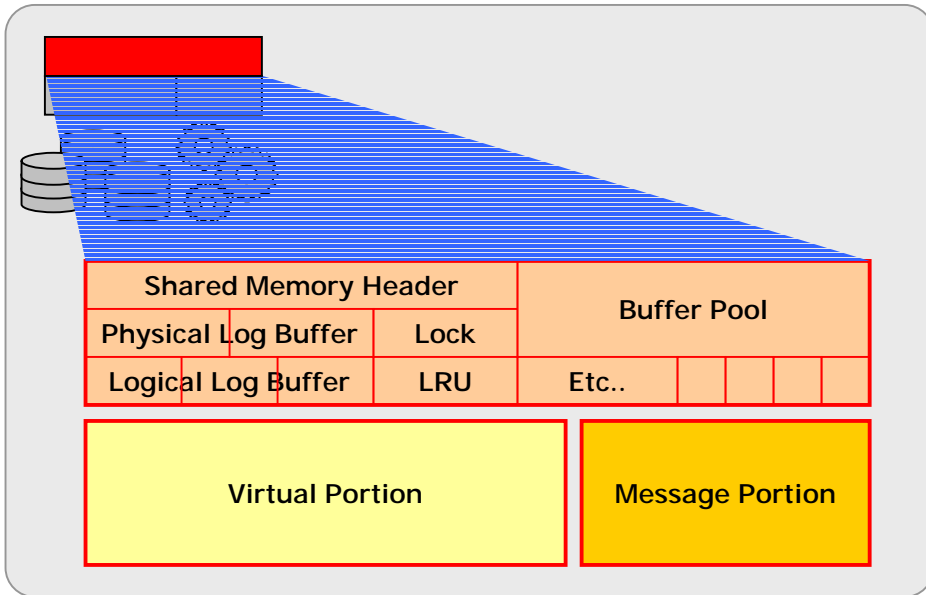
Individual virtual processors:

vp	pid	class	usercpu	syscpu	total
1	21781	cpu	2.84	0.69	3.53
2	21782	adm	0.00	0.03	0.03
3	21783	lio	0.00	0.01	0.01
4	21784	pio	0.01	0.00	0.01
5	21785	aio	0.00	0.02	0.02
6	21786	msc	0.02	0.00	0.02
7	21787	aio	0.00	0.03	0.03
8	21788	shm	0.00	0.01	0.01
9	21789	aio	0.01	0.02	0.03
10	21790	aio	0.01	0.01	0.02
11	21791	aio	0.00	0.01	0.01
12	21792	aio	0.00	0.01	0.01
		tot	2.89	0.84	3.73

```
Informix> vi $INFORMIXDIR/etc/$ONCONFIG
```

```
NETTYPE      SHM,1,10,CPU # Configure poll thread(s) for nettype
NETTYPE      TLI,1,20,NET # Configure poll thread(s) for nettype
NUMCPUVPS    1             # Number of user (cpu) VPs
NUMAIOVPS    3             # Number of IO vps
```


Shared Memory Component : Resident Portion



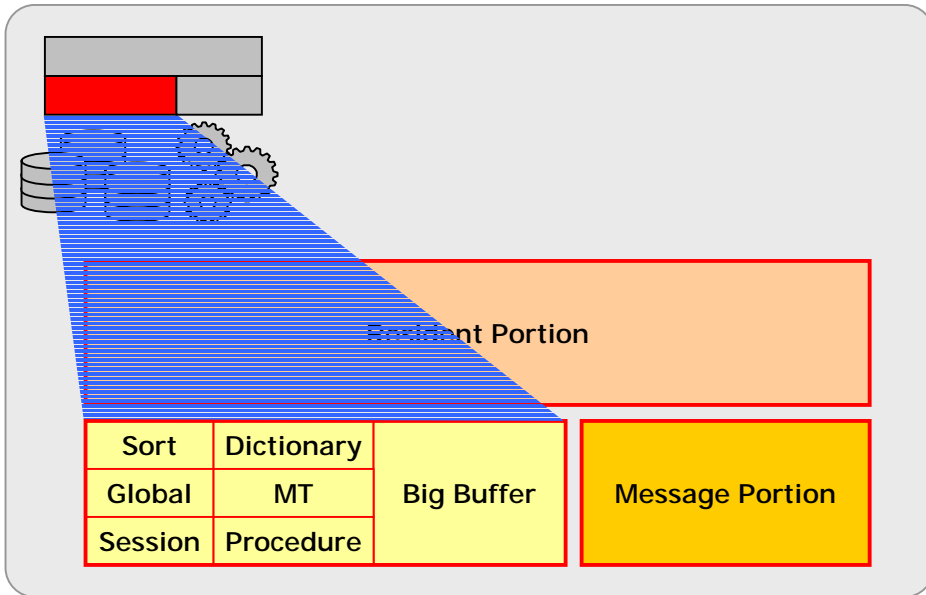
```
Informix> vi $INFORMIXDIR/etc/$ONCONFIG

LOCKS          2000      # Maximum number of locks
BUFFERS         1000      # Maximum number of shared buffers
PHYSBUFF        32       # Physical log buffer size (Kbytes)
LOGBUFF         32       # Logical log buffer size (Kbytes)
LRUS            8        # Number of LRU queues
LRU_MAX_DIRTY  60        # LRU percent dirty begin cleaning limit
LRU_MIN_DIRTY  50        # LRU percent dirty end cleaning limit
```

- **Buffer Pool**
 - Resident Portion 가
 - 가
- **Physical Log Buffer**
 - physical log 가
 - physical log (before image) 가
- **Logical Log Buffer**
 - logical log 3 가
 - logical log DDL, DML 가
- **LRU Queues**
 - LRU(Least Recently Used) linked list pointing dirty
 - FLRU(Free LRU) 가
 - LRU 가
 - User 가 user MLRU 가
 - FLRU 가



Shared Memory Component : Virtual Portion



- shared memory pool 50
- 가 ,
- 8K
- virtual portion
 - private
 - Dictionary :
 - Stored Procedure : MT pool
 - Thread : MT pool
 - Sorting :
 - Big Buffer : AIO VP
 - big buffer가 , 8
 - Global Information :
- Virtual Portion (SHMVIRTSIZE) 가 (SHMADD)
 - config , 가 virtual portion 가
 - release .

```
Informix> vi $INFORMIXDIR/etc/$ONCONFIG
```

```
SHMVIRTSIZE 8000 # initial virtual shared memory segment size
SHMADD 8192 # Size of new shared memory segments (Kbytes)
SHMTOTAL 0 # Total shared memory (Kbytes). 0=>unlimited
```


```
Informix> onmode -a 1
Informix> onstat -g seg
```

Segment Summary:

id	key	addr	size	ovhd	class	blkused	blkfree
2162	1388857345	a000000	2097152	215344	R	274	238
2163	1388857346	a200000	8388608	856	V	1967	81
373	1388857347	aa00000	1048576	632	M	133	123
374	1388857348	ab00000	8388608	856	V	1	2047
652	1388857349	b300000	1048576	632	V	1	255

```
Informix> onmode -F
```

Shared Memory Component : Message Portion



-
-

DBMS

가

가 attach
12Kbytes

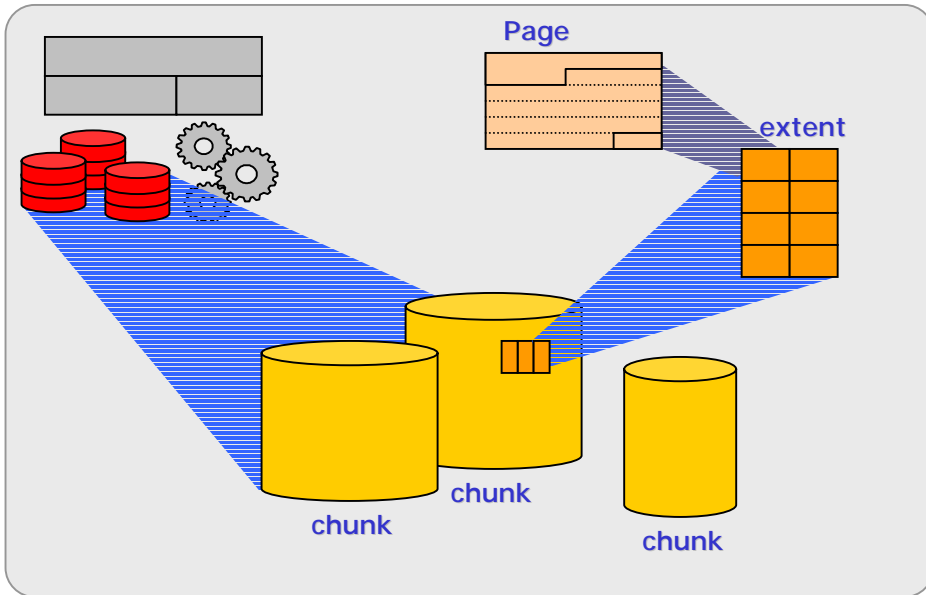
```
Informix> vi $INFORMIXDIR/etc/$ONCONFIG
```

```
NETTYPE      SHM,1,10,CPU # Configure poll thread(s) for nettype
```



```
Informix> vi $INFORMIXSQLHOSTS
```

```
Server_name  onipcshm      host_name     service_port
```

Disk Component : Physical Concept



Raw device Path=/dev/xyz01 Offset=0 Kbytes Size=20000 Kbytes	Piece of Raw device Path=/dev/xyz02 Offset=5000 Kbytes Size=10000 Kbytes	Cooked device Path=/data/chunk1 Offset=0 Kbytes Size=5000 Kbytes
--	--	--



• Chunk

```

-          ,          offset
- chunk mount          character special device raw disk
- offset          raw disk          , mount
-          chunk          2047
- Cooked file          offset 0
- Permission 660 (rw-rw----), informix owner, group
    
```

• Page

```

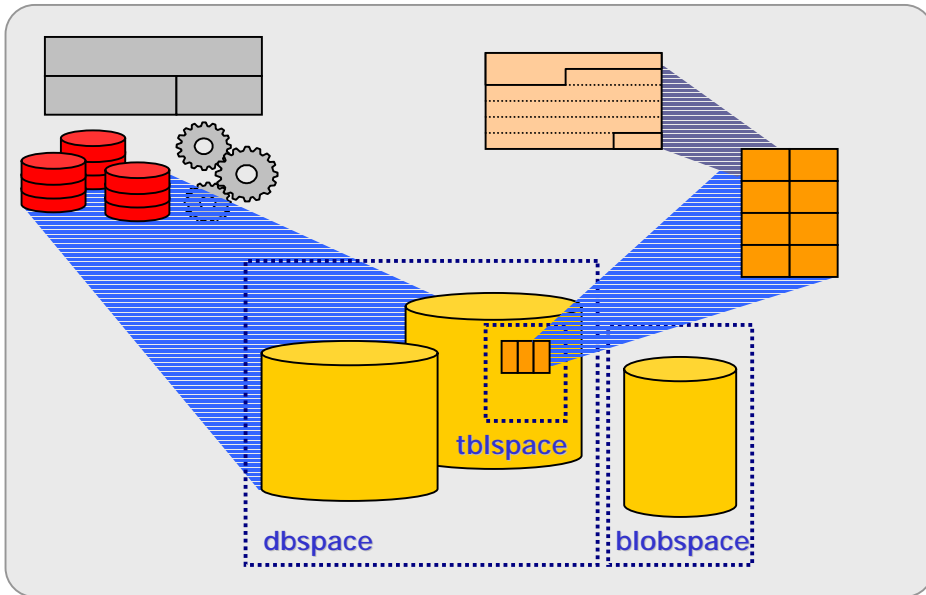
-          I/O          O/S
-          4K          page          2K
-          ,          가          , 4          timestamp, row
-          가 slot
-          timestamp
-          slot(4) : row offset(2) + row size(2)
-          byte          row 4bytes slot          - 24 byte - 4
    
```

• Extent

```

-          . extent
-          EXTENT SIZE :
-          NEXT SIZE : 가
- extent          4          8
-          , 2KB          16KB, 4KB
-          32K가
    
```


Disk Component : Logical Concept



• **Tbldspace**

-
- , 가
- **DBspace**
 - chunk DBspace chunk
 - chunk primary chunk DBspace chunk
 - chunk 가 DBspace
 - DBspace가 , root dbspace
 - 2047 dbspace 가 . 2K
 - root dbspace 40 dbspace

• **Blobspace**

- Simple Large Object (TEXT, BYTE)
- dbspace. blobspace 가 dbspace 2K 4K
- user가 , blobpace
- user가 Large Object
- I/O
- Blob page 가
- logical log large object가
- 가

• **Sbldspace**

- Smart Large Object (BLOB, CLOB), User-defined Data type
- dbspace. sbldspace chunk chunk header data,
- metadata, user data . Meta data user data
- handle random access API
- meta data sbldspace
- . meta data page
- oncheck -ps . header data meta data
- user data .

Informix> onstat -d

Dbspaces							
address	number	flags	fchunk	nchunks	flags	owner	name
a30b7d0	1	0x1	1	1	N	informix	rootdbs
a34a488	2	0x8001	2	1	N S	informix	sbldspace
a98fc68	3	0x2001	3	1	N T	informix	tempdbs
a9a02a0	4	0x11	4	1	N B	informix	blobdbs
4 active, 2047 maximum							

Chunks							
address	chk/dbs	offset	size	free	bpages	flags	pathname
a30b918	1	0	10000	537		PO-	/CHUNK/root.1
a34a320	2	0	500	336	347	POS	/CHUNK/sbld.1
			Metadata	100	74	100	
a98fdb0	3	0	500	447		PO-	/CHUNK/temp.1
a9a03e8	4	0	500	497	500	POB	/CHUNK/blob.1

Informix> vi \$INFORMIXDIR/etc/\$ONCONFIG

```

ROOTNAME      rootdbs      # Root dbspace name
ROOTPATH      /CHUNK/root.1 # Path for device containing root dbspace
ROOTOFFSET    0           # Offset of root dbspace into device (Kbytes)
ROOTSIZE      20000        # Size of root dbspace (Kbytes)
    
```

2. SERVER CONFIGURATION

- Environment Variables
- Disk Parameters
- Shared Memory Parameters
- Performance Tuning Parameters
- Additional Parameters
- How the Client Connects
- Options in the sqlhosts file

Environment Variables

- **INFORMIXDIR**
- **INFORMIXSERVER**
- **PATH**
- **가**
- **ONCONFIG**
- **INFORMIXSQLHOSTS**
- **Code Set**
- **DB_LOCALE, SERVER_LOCALE, CLIENT_LOCALE**
- **Shared Library Path**
- **LD_LIBRARY_PATH, SHLIB_PATH, LIBPATH**
- **PDQPRIORITY**
- **PSORT_NPROCS**
- **OPTCOMPIND**
- **FET_BUF_SIZE**
- **DBSPACETEMP, INFORMIXSTACKSIZE**

- - **INFORMIXDIR** :
 - **PATH** : \$INFORMIXDIR/bin PATH 가
 - **INFORMIXSERVER** : . Configuration DBSERVERNAME DBSERVERALIASES
- 가
 - **ONCONFIG** : Configuration default onconfig가 \$INFORMIXDIR/etc
 - **INFORMIXSQLHOSTS** : sqlhosts \$INFORMIXDIR/etc/sqlhosts 가
- KS5601 sort ,
 - **DB_LOCALE, SERVER_LOCALE, CLIENT_LOCALE** : ko_kr.ksc
 - default en_us.8859-1
 - **SERVER_LOCALE** , DB_LOCALE
- Shared Library path
 - .so .sl
 - \$INFORMIXDIR/lib:\$INFORMIXDIR/lib/esql
 - **LD_LIBRARY_PATH**(SUN,LINUX),**SHLIB_PATH** (HP),**LIBPATH**(AIX)
- - **PDQPRIORITY** : PDQ query PDQ resource %
 - 0 : PDQ query
 - 1 ~ 100 : PDQ resource %
 - **PSORT_NPROCS** : sort sort thread
 - **OPTCOMPIND** : Optimizer Hint
 - 0 : scan
 - 1 : Isolation level repeatable read 0, 2
 - 2 : Optimizer가 query path (default)
 - **FET_BUF_SIZE** : Fetch Buffer size, 가
 - default 4096 bytes 32767 bytes

Disk Parameters

- MIRROR 0 # Mirroring flag (Yes = 1, No = 0)
- TAPEDEV /dev/null
- TAPEBLK 16 # Kbytes
- TAPESIZE 10240 # Kbytes
- LTAPEDEV /dev/null
- LTAPEBLK 16 # Kbytes
- LTAPESIZE 10240 # Kbytes
- ROOTNAME rootdbs
- ROOTPATH /CHUNK/root.1
- ROOTOFFSET 0 # Kbytes
- ROOTSIZE 30000 # Kbytes
- MIRRORPATH
- MIRROROFFSET 0 # Kbytes
- PHYSDBS rootdbs
- PHYSFILE 2000 # Kbytes
- LOGFILES 6
- LOGSIZE 2000 # Kbytes

```
Informix> onstat -l

Physical Logging
Buffer bufused bufsize numpages numwrts pages/io
P-1 0 16 0 0 0:00
phybegin physize phypos phyused %used
100107 1000 589 0 0.00

Logical Logging
Buffer bufused bufsize numrecs numpages numwrts recs/pages pages/io
L-1 0 16 12 6 6 2.0 1.0
Subsystem numrecs Log Space used
OLDRSAM 12 1564
address number flags unqid begin size used %used
a07c680 1 U-B---- 1086 1004ef 250 250 100.00
: : : : : : :
a07c70c 6 U-B---- 1091 1009d1 250 250 100.00
a07c728 7 U--C-L 1092 100acb 250 115 46.00


Informix> onstat -m
14:07:39 Checkpoint Completed: duration was 0 seconds.
14:07:39 Checkpoint logunqid 1092, logpos 0x6e2a8
14:08:08 Checkpoint Completed: duration was 0 seconds.
14:08:08 Checkpoint logunqid 1092, logpos 0x702ac
```

- MIRROR 0 mirroring
- TAPEDEV LTAPEDEV TAPESIZE LTAPESIZE
 - 가 10G
 - , TAPESIZE 1000000 1G
 - 가 TAPE ontape
 - onbar 가 I/O 가
 - TBPEBLK,LTAPEBLK LTAPEDEV가 /dev/null 가 가
 - configuration , LTAPEDEV /dev/null
- ROOTNAME, ROOTPATH, ROOTOFFSET
 - chunk
 - dbspace logical log, physical log
 - 가 가
 - MIRRORPATH, MIRROROFFSET
 - MIRROR가 1 Mirror chunk
 - primary chunk
 - 가
- PHYSDBS, PHYSFILE physical log
 - , LOGFILES,
 - LOGSIZE logical log
 - rootdbs dbspace 가
 - logical log physical log rootdbs
 - dbspace logical log
 - Logical log Physical log
 - 1:2 1:3 , Physical log 가 check
 - point가
 - Physical log가 75% check point interval check
 - point가 flush 가
 - holding



Shared Memory Parameters

- **SERVERNUM** 6
- **DBSERVERNAME** training
- **DBSERVERALIASES** training_tcp
- **DBSPACETEMP** tempdbs1,tempdbs2
- **DEADLOCK_TIMEOUT** 60 # seconds
- **TXTIMEOUT** 300 # seconds
- **HETERO_COMMIT** 0
- **STACKSIZE** 32 # Kbytes



- **SERVERNUM** segment 가
- **DBSERVERNAME** \$INFORMIXSQLHOSTS , INFORMIXSERVER가
- **DBSPACETEMP** order by, group by, select .. into temp.. SQL temporary dbspace dbspace I/O가 **DBSPACETEMP** 가

SERVERNUM 0 255

SERVERNUM : 0					SERVERNUM : 1					SERVERNUM : 6				
0x	52	56	48	01	0x	52	57	48	01	0x	52	5C	48	01
0x	52	56	48	02	0x	52	57	48	02	0x	52	5C	48	02
0x	52	56	48	03	0x	52	57	48	03	0x	52	5C	48	03

```
Informix> ipcs -m
m      2162 0x525c4801 -rw-rw--- root informix
m      2163 0x525c4802 -rw-rw--- root informix
m      373 0x525c4803 -rw-rw-rw- root informix
```


```
Informix> onstat -g seg
id key addr size ovhd class blkused blkfree
2162 1381779457 aa00000 2097152 215344 R 274 238
2163 1381779458 aa200000 8388608 856 V 1973 75
373 1381779459 aa00000 1048576 632 M 133 123
```

```
Informix> vi $INFORMIXSQLHOSTS
training onipcshm host1 informix_shm
training_tcp onlitcp host1 service_port1
```

- **DEADLOCK_TIMEOUT** dead lock
- **TXTIMEOUT** 2 phase-commit 가
- **HETERO_COMMIT** 1 , 2 phase commit 0
- **STACKSIZE** thread가 configuration (default 32KB), User session **INFORMIXSTACKSIZE** stack

Shared Memory Parameters

- **CLEANERS** 2
- **SHMTOTAL** 0 # Kbytes. 0=>unlimited
- **SHMVIRTSIZE** 8000 # Kbytes
- **SHMADD** 8192 # Kbytes
- **PHYSBUFF** 32 # Kbytes
- **LOGBUFF** 32 # Kbytes
- **LOCKS** 2000
- **BUFFERS** 5000



```

Informix> onstat -g ath
13 a85ddb0 a30c608 2 sleeping forever 1cpu flush_sub(0)
14 a8f5018 a30cbf8 2 sleeping forever 1cpu flush_sub(1)

Informix> onstat -F
address flusher state data
a30c608 0 I 0 = 0X0
a30cbf8 1 I 0 = 0X0
  
```

```

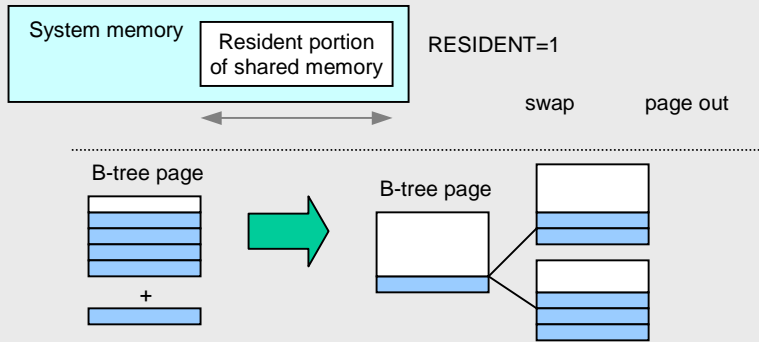
Informix> onstat -p
dskreads pagreads bufreads %cached dskwrits pagwrits bufwrits %cached
126 162 3354 96.24 7 7 3 0.00

ovlock ovuserthread ovbuff usercpu syscpu numckpts flushes
0 0 0 2.88 0.78 1 14
  
```

- Page Cleaner dirty page(MLRU pointing AIO(KAIO) queue) CLEANERS page cleaner
- guideline
 - 20 100 1 2 1
 - 100 4 1
 - Check point duration page cleaner thread 가
- SHMVIRTSIZE virtual portion , SHMADD virtual portion 가 virtual portion , 0 . Virtual segment가 SHMTOTAL shared memory , 0 . Virtual segment가 SHMVIRTSIZE virtual portion 가 segment가 virtual portion 가 segment가 가 : onmode -a seg_size virtual portion : onmode -F
- PHYSBUFF, LOGBUFF physical log, logical log 가 flush I/O가
- LOCKS 가 Lock (9.2 가)
- BUFFERS 2K resident portion buffer 5000 * 2K 가 physical memory 20% . onstat -p read cache 95%

Shared Memory Parameters

- ONDBSPACEDOWN 2 # 0 = CONTINUE, 1 = ABORT, 2 = WAIT
- RESIDENT 1 # Yes = 1, No = 0
- LTXHWM 50 # percentage
- LTXEHWM 60 # percentage
- FILLFACTOR 90 # percentage



```
Informix> onstat -g seg
id      key      addr      size      ovhd      class blkused blkfree
3062    1388857345 a000000    16777216  215792    R*      274      3822

Informix> onstat -g tpf
tid      lkreqs  lkwr  dl  to  lgrs  isrd  iswr  isrw  isdl  isct  isrb  ix  bfr  bfw  lsus  lsmx  seq
1068    333    0    0    0    0    136  0    0    0    0    0    1  428  0    0    0    5


Informix> onstat -g ath
tid      tcb      rsrcb      prty status      vp-class      name
1068    d1ec7918 d13e6380  2    cond wait    netnorm      1cpu         sqlexec

Informix> onstat -u
address  flags  sessid  user  tty  wait  tout  locks  nreads  nwrites
d13e6380 Y--P--- 1067920 informix 4    d1d75de8 0    1    62    0
```

- ONDBSPACEDOWN rootdbs log가 dbspace가
 - 0 (continue) : dbspace chunk "D" flag
 - ontape -r -D dbspace dbspace
 - 1 (abort) : 가 . cold restore
 - ontape -r
 - 2 (wait) : checkpoint "D" flag
 - 가
 - checkpoint 가
 - chunk flag "D" .
- shared memory resident portion swap
 - page out RESIDENT 1 . ,
 - 가 .
- LTXHWM
 - long transaction rollback . Rollback
 - 가 LTXEHWM , long transaction
 - rollback suspend . LTXHWM, LTXEHWM
 - 9.3 .
- FILLFACTOR b-tree index %
 - 가 . create index .
 - create index index name on table name(column name) fillfactor 80

Performance Tuning Parameters

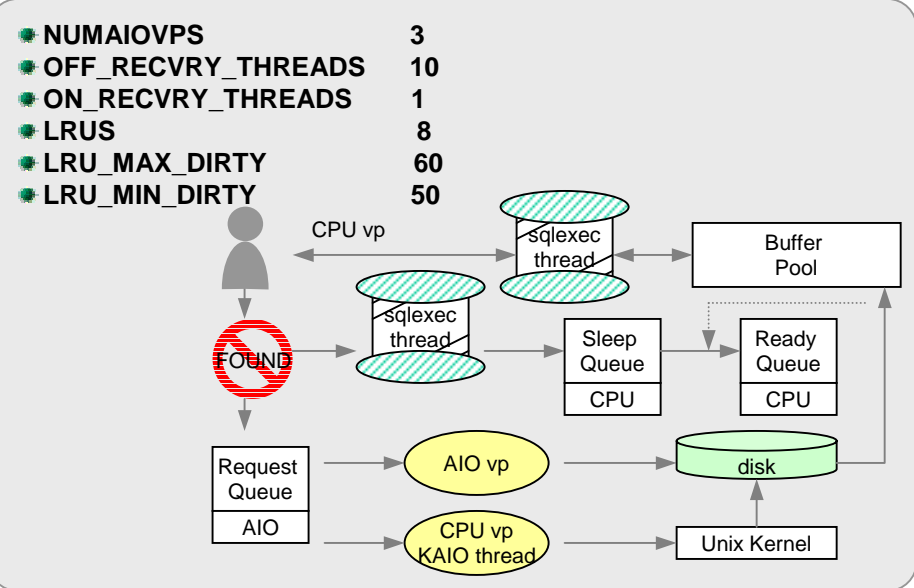
- MULTIPROCESSOR 1
- SINGLE_CPU_VP 0
- AFF_SPROC 1
- AFF_NPROCS 2
- NUMCPUVPS 1
- USEOSTIME 0 # 0: internal (fast), 1: OS(slow)
- NOAGE 0




```
Informix> onstat -g glo
vp  pid  class  usercpu  syscpu  total
1   17272  cpu     2.99    0.64    3.63
Informix> onstat -g ath
7    b677238 b10c018 4   sleeping secs: 1      1cpu   main_loop()
8    b677d18 0       2   running              1cpu   tlitcpoll
10   b6ac768 0       3   sleeping forever     1cpu   tlitcplst
11   b635308 0       3   sleeping forever     1cpu   sm_listen
:
Informix> onmode -p +1 cpu
Informix> onstat -g glo
vp  pid  class  usercpu  syscpu  total
1   17272  cpu     2.99    0.64    3.63
15  17294  cpu     0.00    0.01    0.01
Informix> onstat -g ath
7    b677238 b10c018 4   sleeping secs: 1      15cpu  main_loop()
8    b677d18 0       2   running              1cpu   tlitcpoll
10   b6ac768 0       3   sleeping forever     1cpu   tlitcplst
11   b635308 0       3   sleeping forever     1cpu   sm_listen
:
Informix> onmode -p -1 cpu
```

- CPU 가 2 NUMCPUVP 2
 - MULTIPROCESSOR 1 VP 0 read-ahead
 - default
 - spin lock , mutex
 - lock lock wait 가
 - (spin)
- CPU vp SINGLE_CPU_VP 1
 - CPU vp 0 SINGLE_CPU_VP가 1
 - mutex call locking bypass
- CPU vp가 CPU binding
 - processor logical 0 processor binding processor
 - AFF_SPROC , AFF_NPROCS
 - process affinity
 - \$INFORMIXDIR/release
- NUMCPUVPS CPU CPU vp
 - Client DB 가 CPU - 1
 - DB CPU 1
 - onmode -p +1 cpu
 - onmode -p -1 cpu
- USEOSTIME 1 SQL current datetime
 - mille second 가 select
- O/S CPU
 - (priority) priority aging
 - running CPU vp
 - priority aging disable NOAGE 1

Performance Tuning Parameters



- NUMAIOVP Disk I/O vp . Kernel Asynchronous I/O 2~3 (Raw device) Kernel Asynchronous I/O (Cooked file) Disk I/O channel .




- OFF_RECVRY_THREADS 가 logical log roll forward recovery cold restore . Fast recovery

- ON_RECVRY_THREADS 가 (restore) . Warm restore

- LRUS Least Recently Used Queue . LRU FLRU queue MLRU queue가 , LRU queue 가 가 page cleaner thread 가 LRU queue 500~700

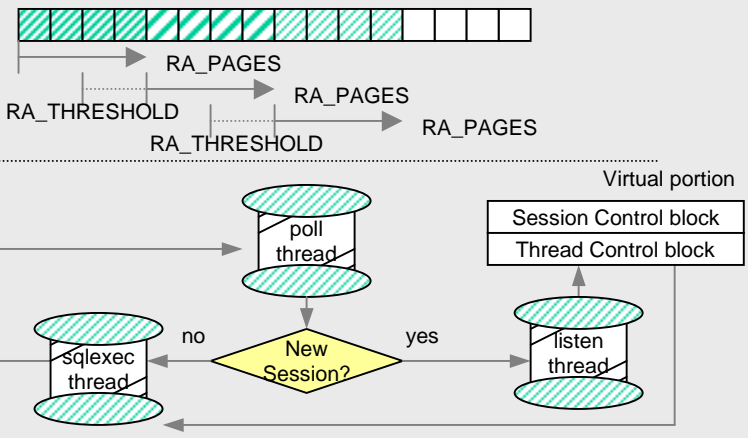
```
Informix> onstat -R
8 buffer LRU queue pairs
#..f./m pair total % of length LOW MED_LOW MED_HIGH HIGH
0 f 26 100.0% 26 22 4 0 0
1 m 0 0.0% 0 0 0 0 0
2 f 26 100.0% 26 22 3 1 0
3 m 0 0.0% 0 0 0 0 0
4 f 26 100.0% 26 23 3 0 0
5 m 0 0.0% 0 0 0 0 0
6 f 26 100.0% 26 24 2 0 0
7 m 0 0.0% 0 0 0 0 0
8 F 24 100.0% 24 20 4 0 0
9 m 0 0.0% 0 0 0 0 0
10 f 24 100.0% 24 21 3 0 0
11 m 0 0.0% 0 0 0 0 0
12 f 24 100.0% 24 20 4 0 0
13 m 0 0.0% 0 0 0 0 0
14 f 24 100.0% 24 20 4 0 0
15 m 0 0.0% 0 0 0 0 0
0 dirty, 200 queued, 200 total, 256 hash buckets, 2048 buffer size
start clean at 60% (of pair total) dirty, or 15 buffs dirty, stop at 50%
0 priority downgrades, 0 priority upgrades
```

- Check point duration check point interval LRU write . LRU write LRU_MAX_DIRTY LRU_MIN_DIRTY . (, LRU_MAX_DIRTY=60, LRU_MIN_DIRTY=40 LRU modified buffer가 60% LRU write가 modified buffer가 40%가 LRU write가)
- LRU write checkpoint가 write checkpoint duration . 가 sync. checkpoint fast recovery . , LRU write



Performance Tuning Parameters

- **CKPTINTVL** 300 # seconds
- **RA_PAGES** 4
- **RA_THRESHOLD** 2
- **NETTYPE** soctcp,1,10,NET



- Check point
 - . Check point가
 - check point duration
 - duration
 - . Check point interval
 - . Check point가
 - CKPTINTVL
 - , physical log가 75%
 - CKPTINTVL
 - Check point가
 - . check point가
 - check
- Read Ahead
 - sequential scan
 - shared memory
 - data page
 - index
 - page sqlexec
 - sqlexec waiting
 - . Read Ahead가
 - page
 - RA_PAGES**
 - . Default 4pages
 - Read Ahead
 - page
 - 0
 - . Read Ahead
 - Read Ahead
 - Read Ahead
 - page
 - memory
 - RA_THRESHOLD**
 - Read Ahead
 - page

```

Informix> onstat -p
ovlock  ovuserthread  ovbuff  usercpu  syscpu  numckpts  flushes
0        0                0        3.01     0.80    1         98

Informix> onstat -g ntu
netscb type  thread name  sid  fd  poll  reads  writes q-nrm q-pvt q-exp
b6d3910 tlitcp  sqlexec    16  3   5     7     7  0/ 1  1/ 1  0/ 0
b6bbb40 ipcshm  sm_discon  8    0   0     0     0  0/ 0  0/ 0  0/ 0
b6c24c8 ipcshm  sm_listen  5    0   0     0     0  0/ 0  0/ 0  0/ 0
b69ccd8 tlitcp  tlitcplst  4    1   5     2     0  0/ 0  0/ 0  0/ 0
b69c400 ipcshm  sm_poll    3    0   0     0     0  0/ 0  0/ 0  0/ 0
b5564a0 tlitcp  tlitcpoll  2    0   5     7     0  0/ 0  0/ 0  0/ 0

Informix> onstat- g ntt
Individual thread network information (times):
netscb thread name  sid  open  read  write address
b6d3910 sqlexec    16  19:06:03  19:06:11  19:06:11
b6bbb40 sm_discon  8   14:28:45
b6c24c8 sm_listen  5   14:28:45
b69ccd8 tlitcplst  4   14:28:45  19:06:03  kor-olive|23456|tlitcp
b69c400 sm_poll    3   14:28:45
b5564a0 tlitcpoll  2   19:06:01
    
```

- **NETTYPE** Client message sqlexec poll thread
- **NETTYPE** protocol, Threads, Users, Vp-class
 - 1st position : poll thread가 protocol connection type
 - poll thread가 , sqlhosts protocol
 - (ipcshm,ipcstr,tlitcp,tlistpx,soctcp)
 - 2nd position : protocol poll thread listen
 - thread 가
 - poll thread 가 . listen thread
 - DBSERVERNAME** , **DBSERVERALIASES**
 - listen thread가
 - 3rd position : poll thread가 message
 - Shared memory connection
 - poll thread 100 user가
 - 4th position : poll thread가 vp class (NET/CPU),
 - CPU vp class 가 NET vp class
 - 가 CPU vp poll thread

Additional Parameters

```

• MSGPATH          /usr/informix/online.log
• CONSOLE          /dev/console
• ALARMPROGRAM    /usr/informix/etc/log_full.sh

• DUMPSHMEM       1
• DUMPPDIR        /tmp
• DUMPCNT         1

• VPCLASS         cpu, num=8, aff=0-7, noage

```

```

• MSGPATH
• CONSOLE          /dev/console
                  /dev/null
  (NT NUL)
• ALARMPROGRAM
  , log_full.sh
  log backup
  log full
• DUMPSHMEM 1
  shmem.uniqid
  dump
  assertion failure가
  shared memory가 dump
  shared memory
  shared memory
• Shared Memory가 dump
  DUMPPDIR (default
  /tmp)
• DUMPCNT
  assertion failure
  shared memory가 dump
  , 1
  dump
  assertion failure가
  shared memory
  가
  VP class
• VPCLASS
  . III-behaved UDR(User Defined Routine)
  UDR CPU VP가
  user-defined VP
  . (UDR
  running VP class
  )
  - ) VPCLASS myvp,num=2,noage ==> myvp
  VP class 2
  - , user defined VP
  , SINGLE_CPU_VP가 0
• JVP vp class CPU VP JVM(Java Virtual Machine) code imbedded
  , java application CPU VP JVM shipping cost


• VPCLASS
  NUMCPUVPS, NUMAIOVPS
  , cpu vp 2 , Processor affinity 2 cpu
  3
  , noage
  - 1) NUMCPUVPS 2
    AFF_SPROC 2
    AFF_NPROCS 3
    NOAGE 1
  - 2) VPCLASS cpu,num=2,aff=2-4,noage
  - 1 2
  1

```

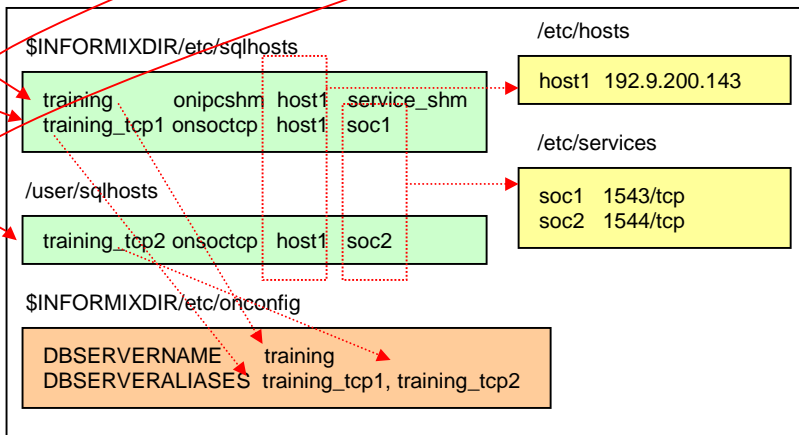


How the Client Connects

- 1. \$INFORMIXSERVER environment variable is used
- 2. Read the \$INFORMIXDIR/etc/sqlhosts file
- 3. Read the DBSERVERNAME, DBSERVERALIASES parameter in the onconfig file

sqlhosts entry



NETTYPE field



- NETTYPE connection H/W , SUN
TLI , HP SOC connection
, \$INFORMIXDIR/release IDS
- Shared memory (ipcshm) sqlhosts
SERVICENAME 가 ,
- Sqlhosts HOSTNAME SERVICENAME
host ip, port
- onconfig DBSERVERALIASES ,
service port DBSERVERNAME,
DBSERVERALIASES listen thread가
- Multi-Ethernet card host IP
가 Ethernet card listen thread
DBSERVERALIASES 가 ..
- \$INFORMIXDIR/etc/\$ONCONFIG
DBSERVERALIASES training_tcp1, training_tcp2, training_tcp3
-\$INFORMIXSQLHOSTS
training_tcp1 onsoctcp host1 soc1
training_tcp2 onsoctcp host2 soc2
training_tcp3 onsoctcp host3 soc3
- /etc/hosts, /etc/services NT windows
, windows95,98 c: \ windows , NT windows2000
c: \ WINNT \ system32 \ drivers \ etc hosts, services

Options in the sqlhosts file

- Keep alive option

- security option

- buffer option

DBSERVERNAME	NETTYPE	HOSTNAME	SERVICENAME	OPTIONS
				k=0,s=2,r=0,b=3072

→ k=0,s=2,r=0,b=3072

- network security files

- /etc/hosts.equiv

- .rhosts

- .netrc

- Keep alive option client server가 connection
 - k=1 : enable keep alive option (default, recommends)
 - k=0 : disable keep alive option
- Security System security
 - r=0 : disables .netrc look-up
 - r=1 : enables .netrc look-up (default)
 - s=0 : disables /etc/hosts.equiv, .rhosts look-up
 - s=1 : enables /etc/hosts.equiv look-up
 - s=2 : enables .rhosts look-up
 - s=3 : enables /etc/hosts.equiv, .rhosts look-up (default)
- Buffer TCP/IP row BLOB client server
- Password 가 trust /etc/hosts.equiv . User user .rhosts . rlogin , password
- .netrc machine ID, user ID, password ,

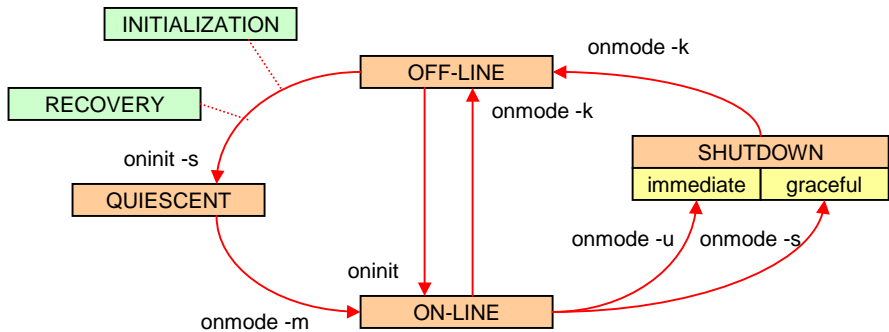
3. SERVER OPERATION

Operating Modes
Shared Memory Buffer Access
Checkpoints
Buffer Writes
Fast Recovery
Step to Fast Recovery Process
Media Failures

Operating Modes

- Off-Line
 - Quiescent
 - On-Line
 - Shutdown (Immediate/ Graceful)
 - Initialization
 - Recovery
-
- oninit [-s]
 - onmode [-m] [-s] [-u] [-y]

- **Off-Line** : 가 . Shared memory가
- **Quiescent** : oninit process가 , shared memory 가 , informix 가
- **On-Line** : 가 가
- **Shutdown** : Quiescent , On-Line . Immediate shutdown , Graceful shutdown 가
- **Initialization** : Off-Line Quiescent
- **Recovery** : 가 fast recovery , restore Off-Line Quiescent
- Shared memory base key 0x52564801 SERVERNUM 0x10000
 - Key = 0x52564801 + (SERVERNUM * 0x10000)



```

Informix> onstat -c | grep SERVERNUM
SERVERNUM      6
Informix> bc
Obase=16
Ibase=16
52564801 + (6 * 10000)
525C4801

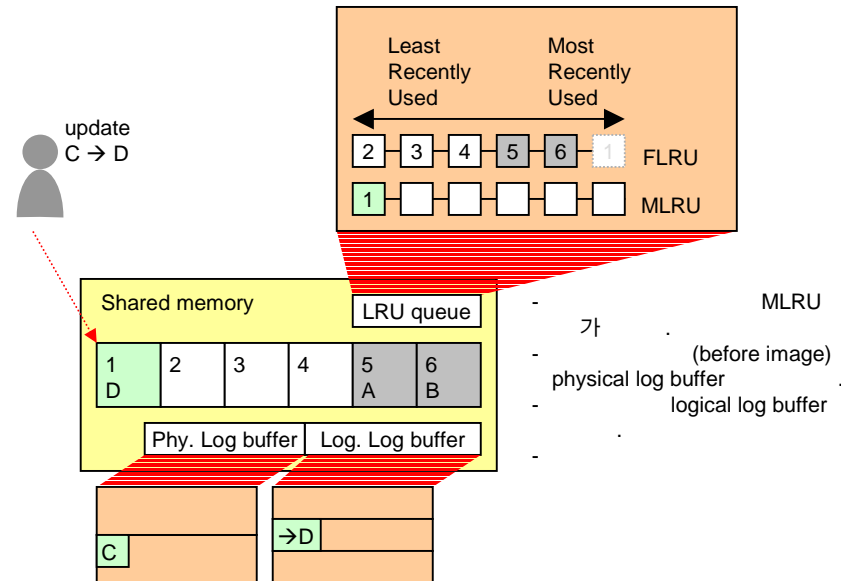
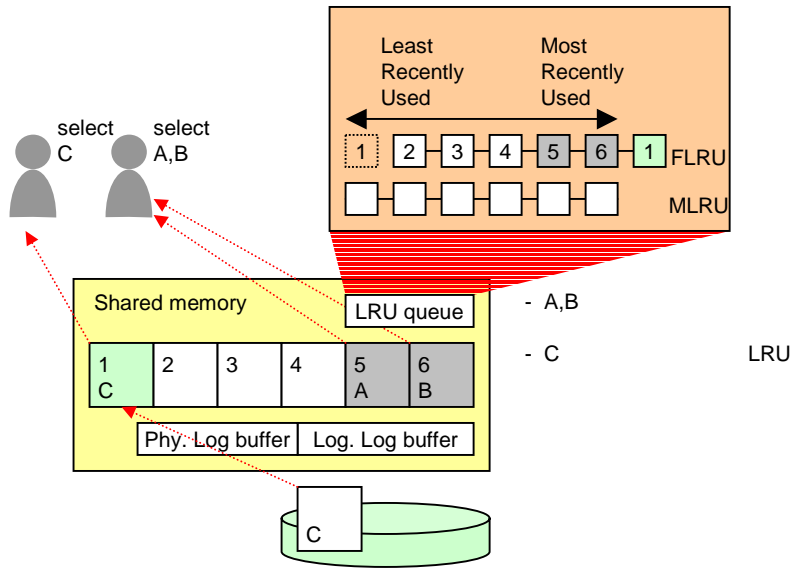
Informix> ipcsc -m
m   541 0x525c4801 --rw-rw---- root informix
m   542 0x525c4802 --rw-rw---- root informix
    
```

Shared Memory Buffer Access

- 1. shared memory buffer most recently used
- 2. LRU 가
- 3. least recently used 가
- 4. LRU 가
- 5. FLRU가 MLRU
- 6. physical log buffer
- 7. logical log buffer

- shared memory mutex, shared memory lock, mutex, document
- User가 가 LRU queue 가
- shared lock, exclusive lock concurrency data consistency

3



Checkpoints

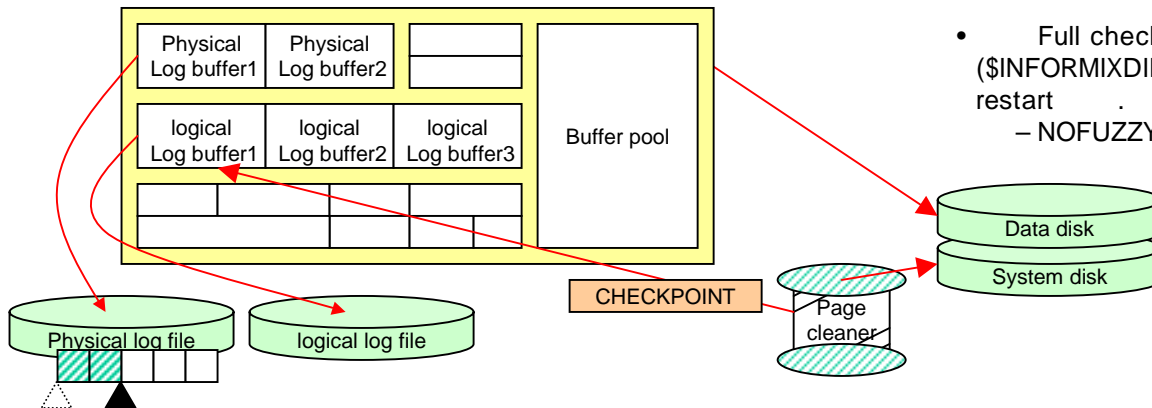

checkpoint shared memory buffer

- Full checkpoint :
- Fuzzy checkpoint : fuzzy operation
 - fuzzy operation – insert, update, delete

checkpoint

- Step 1 : User thread critical section
- Step 2 : Physical Log buffer가 flush
- Step 3 : buffer pool page가 flush
(Fuzzy checkpoint fuzzy operation)
- Step 4 : Logical log buffer checkpoint-complete
System reserved page checkpoint
- Step 5 : Logical Log buffer가 flush
- Step 6 : Physical Log

- Checkpoint duration user blocking 9.2 full
- Fuzzy checkpoint modified buffer가 LRU write가
- Fuzzy checkpoint modified buffer가 LRU write가
- Fuzzy checkpoint
 - Checkpoint interval (CKPTINTVL) 가
 - Physical Log가 75%
 - logical log checkpoint가
 - Chunk dbspace
 - `onmode -c fuzzy`
- Full checkpoint
 - Ontape onbar
 - Fast recovery (oldest) checkpoint가 logical log
 - `onmode -c`
 - `onmode -yuck`



- Full checkpoint (\$INFORMIXDIR/etc/\$ONCONFIG) restart
 - NOFUZZYCKPT 1
- configuration file 가 DB

Buffer Writes

Buffer write

- Chunk Write
- LRU Write
- Foreground Write

Flushing the Log buffer

- Physical log buffer
 - checkpoint
 - physical log buffer flush
- Logical Log buffer
 - checkpoint
 - logical log buffer flush

```
Informix> onstat -F
Fg Writes    LRU Writes    Chunk Writes
0            0             37
```

```
address  flusher  state  data
b10c608  0        I      0      = 0X0
b10cbf8  1        I      0      = 0X0
states: Exit Idle Chunk Lru
```

• Chunk write

- Checkpoint write
- Page cleaner가 modified buffer AIO chunk write, page sort write, seek time

• LRU write

- LRU_MAX_DIRTY LRU 가 modify flush, chunk write user thread blocking, sort, LRU write가, checkpoint가, chunk write, user thread blocking, checkpoint, duration, sort, write, chunk write가

• Foreground write


- LRU 가 free buffer FLRU access, MLRU 가 write foreground write, write page cleaner가 user thread가 LRU write가

• Physical log buffer가 가

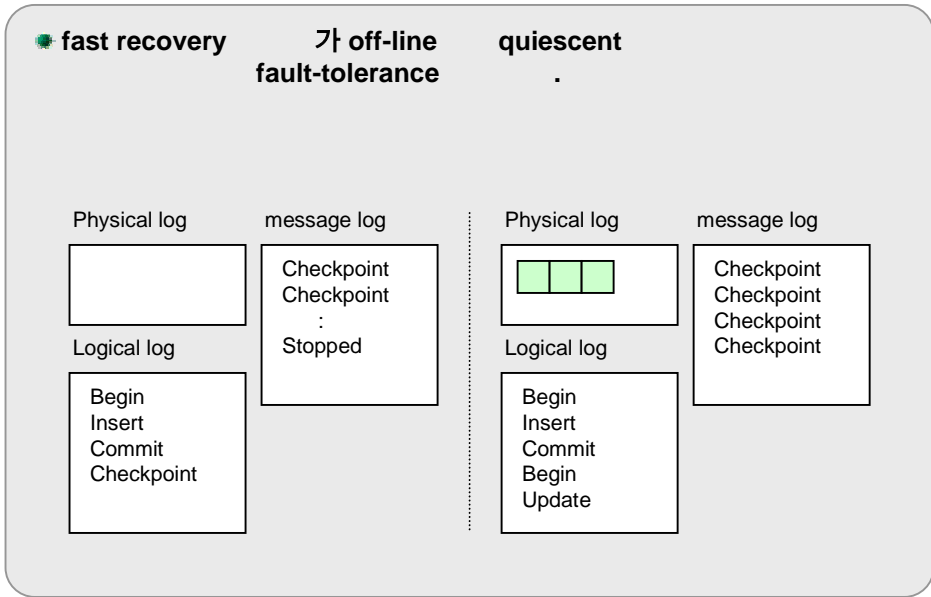
- checkpoint가 physical log buffer flush, physical log buffer가 flush, physical log buffer flush

• checkpoint가

- logical log buffer가 flush, logging Logical Log buffer가 flush
- Buffered logging : logical log buffer가 가 flush
- Un-buffered logging : logical log buffer flush



Fast Recovery



- fast recovery DB H/w crash power , DB fast recovery 가
- fast recovery checkpoint physical log logical log checkpoint roll-forward roll-back
- shutdown checkpoint가 physical log , logical log checkpoint record message log shutdown 가
- 가 checkpoint가 가 physical log logical log physical log가 checkpoint record가 fast recovery가


Step to Fast Recovery Process

- Step 1 : physical log before image
- Step 2 : logical log fast recovery
 - Full checkpoint – checkpoint
 - Fuzzy checkpoint – 가 flush
- Step 3 (Fuzzy checkpoint only) : checkpoint flush
- Step 4 : checkpoint logical log operation
- Step 5 : (roll forward) operation
- Step 5 : roll back

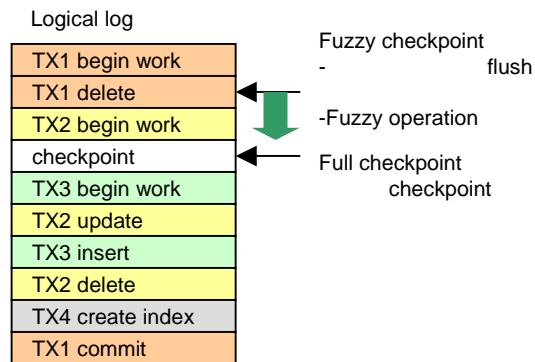
- 가 LRU write flush
- fast recovery physical log checkpoint logical log
- Fuzzy checkpoint fuzzy operation(update, delete, insert) flush

```
Informix> onstat -m
Mon Feb 14 16:18:18 1994
16:18:18 INFORMIX-OnLine Initialized -- Shared Memory Initialized
16:18:18 Physical Recovery Started
16:18:18 Physical Recovery Complete: 5 Pages Restored
16:18:18 Logical Recovery Started
16:18:18 Logical Recovery allocating 3 worker threads ('OFF_RECVR_THREADS').
16:18:22 Logical Recovery Complete
        11 Committed, 2 Rolled Back, 0 Open, 0 Bad Locks
16:18:24 Quiescent Mode
16:18:25 Checkpoint Completed: duration was 1 seconds
16:18:37 On-Line Mode
```

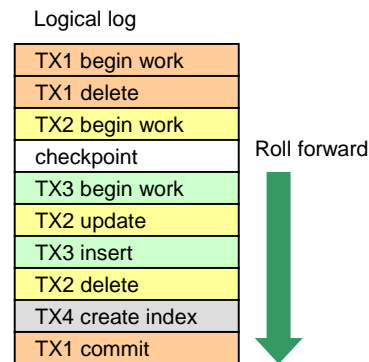
Step 1



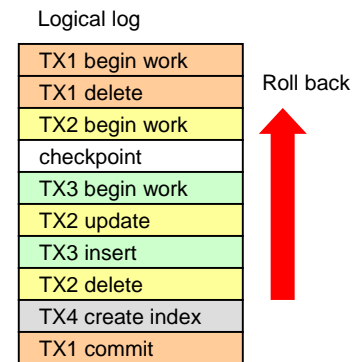
Step 2, 3



Step 4




Step 5




Media Failures

- Mirroring configuration**
 - primary chunk I/O 가 ,
 - mirror chunk switch .
 - ONDBSPACEDOWN
 - I/O가
 - 0 = continue (rootdbs log dbSPACE가)
 - 1 = abort
 - 2 = wait (default. checkpoint user thread가 hang)


- Mirroring dbSPACE user**
 - 가 가




```
Informix> onstat -d
Chunks
address  chk/dbs offset  size  free  bpages  flags  pathname
b10b918  1  1  0      10000  0      PD-   /CHUNK/root.1
b10ba80  1  1  0      10000  537    MO-   /CHUNK/mroot.1
```



```
Informix> onspaces -s rootdbs -p /CHUNK/root.1 -o 0 -0 -y
Chunk status successfully changed.
Informix> onstat -d
Chunks
address  chk/dbs offset  size  free  bpages  flags  pathname
b10b918  1  1  0      10000  0      PR-   /CHUNK/root.1
b10ba80  1  1  0      10000  537    MO-   /CHUNK/mroot.1
```



```
Informix> onstat -d
Chunks
address  chk/dbs offset  size  free  bpages  flags  pathname
b10b918  1  1  0      10000  0      PO-   /CHUNK/root.1
b10ba80  1  1  0      10000  537    MO-   /CHUNK/mroot.1
```



```
Informix> onspaces -m rootdbs -p /CHUNK/root.1 -o 0 -m /CHUNK/mroot.1 0 -y
A system archive should be created after mirroring a chunk that contains logs.
Verifying physical disk space, please wait ...
The Space "rootdbs" is now mirrored.
```

```
Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b10b7d0 1 0x2 1 1 M informix rootdbs
Chunks
address  chk/dbs offset  size  free  bpages  flags  pathname
b10b918  1  1  0      10000  0      PO-   /CHUNK/root.1
b10ba80  1  1  0      10000  537    MX-   /CHUNK/mroot.1
```

```
Informix> ontape -s
```

```
Informix> onstat -d
Chunks
address  chk/dbs offset  size  free  bpages  flags  pathname
b10b918  1  1  0      10000  0      PO-   /CHUNK/root.1
b10ba80  1  1  0      10000  537    MO-   /CHUNK/mroot.1
```

4. *DISK MANAGEMENT*

Disk Space for a Table

Chunk Layout

Space Management

Physical Log / Logical Log

Log Management

Disk Space for a Table

review

- **Chunk :**
- **Page :**
- **Extent :**
- **tblspace :** extent


extent

- **Bitmap Page :** extent
- **Data Page :**
- **Remainder Page :** page
- **Blob Page :** blob
- **Free Page :** extent
- **Index page :** index

extent growth

- **concatenation –** extent extent 가
- **doubling – 16** extent 2
- **manual modification – alter table** next extent

Concatenation



Manual

Alter table table_name modify next size 128


Table extent

Page 0 bitmap	Page 1 data
Page 2 data	Page 3 data
Page 4 remainder	Page 5 blob
Page 6 free	Page 7 free

index extent

Page 0 bitmap	Page 1 index
Page 2 index	Page 3 index
Page 4 index	Page 5 free
Page 6 free	Page 7 free

Doubling




- Extent tblspace 8
- Extent가 200 re-org. extent 가
- 가 extent
- tblspace가
- Bitmap page tracking 4bit extent
- Data page slot integer 4byte rowid Slot rowid page unsigned 1 byte 255
- Remainder page Home remainder 4byte 가 가 dbspace blob
- Blob space blob
- Index page (9.2 delete_flag) Index page row key value 4byte rowid, 1byte index 가 extent 가 b-tree level

Exercise 1

1. DB, TABLE

- create database new_db;
- create table new_table (col1 char(800));
- oncheck -pT new_db:new_table
- 1 row
- insert into new_table values ('1st row');
- oncheck -pT new_db:new_table
- 3,4
- 6.
- alter table new_table add (col2 char(1600));
- update new_table set col2 = 'hello';
- oncheck -pT new_db:new_table

```
Informix> oncheck -pT new_db:new_table
Type          Pages    Empty  Semi-Full    Full  Very-Full
-----
Free          7
Bit-Map       1
Index         0
Data (Home)   0
-----
Total Pages   8
Unused Space Summary
Unused data slots          0
Unused bytes per data page 412
Total unused bytes in data pages 0
```




extent

Page header : 24
 Time stamp : 4
 2 row 가 : 800 * 2 = 1600
 2 row slot : 4 * 2 = 8
 Total : 1636
 Unused byte : 2048-1636 = 412


• 1 row

```
Informix> oncheck -pT new_db:new_table
Type          Pages    Empty  Semi-Full
-----
Free          6
Bit-Map       1
Index         0
Data (Home)   1
-----
Total Pages   8
Unused Space Summary
Unused data slots          1
Unused bytes per data page 412
Total unused bytes in data pages 412
```




• 2 row

```
Informix> oncheck -pT new_db:new_table
Type          Pages    Empty  Semi-Full
-----
Free          5
Bit-Map       1
Index         0
Data (Home)   2
-----
Total Pages   8
Unused Space Summary
Unused data slots          1
Unused bytes per data page 412
Total unused bytes in data pages 824
```



• Schema

```
Informix> oncheck -pT new_db:new_table
Type          Pages    Empty  Semi-Full
-----
Free          3
Bit-Map       1
Index         0
Data (Home)   3
Data (Remainder) 1          0
-----
Total Pages   8
Unused Space Summary
Unused data slots          0
Unused data bytes in Remainder pages 844
```



Exercise 2

1. TABLE, INDEX

- create table new_table2 (col1 char(255));
- create unique index idx_table2 on new_table2 (col1);

2. 1 row

- insert into new_table2 values ('1st row');



3. oncheck -pT new_db:new_table2

4. 7 row

5. oncheck -pT new_db:new_table2



• 1 row

```
Informix> oncheck -pT new_db:new_table2
Type          Pages      Empty  Semi-Fu
-----
Free          6
Bit-Map       1
Index         0
Data (Home)   1
-----
Total Pages   8
Unused Space Summary
  Unused data slots          6
  Unused bytes per data page 207
  Total unused bytes in data pages
    Average  Average
Level  Total No. Keys Free Bytes
-----
1      1      1      1756
-----
Total  1      1      1756
```



• 7 row

```
Informix> oncheck -pT new_db:new_table2
Type          Pages      Empty  Semi-Fu
-----
Free          5
Bit-Map       1
Index         0
Data (Home)   2
-----
Total Pages   8
Unused Space Summary
  Unused data slots          6
  Unused bytes per data page 207
  Total unused bytes in data pages
    Average  Average
Level  Total No. Keys Free Bytes
-----
1      1      2      1749
2      2      4      964
-----
Total  3      3      1225
```

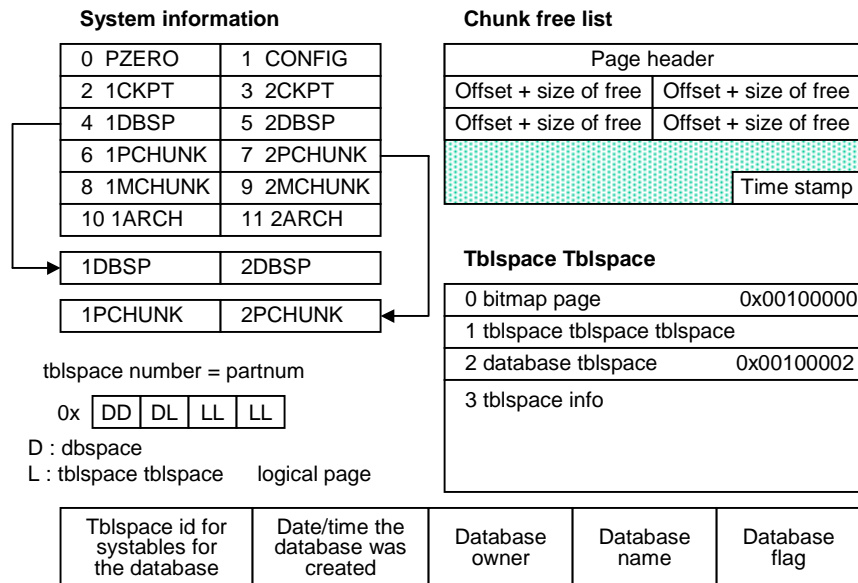


Chunk Layout

- System Overhead
 - system information (Reserved Page)
 - Chunk Free List
 - Tblspace Tblspace
- Physical Log
- Logical Log
- sysmaster database
- syutils database

Rootdbs	chunk	dbspace	chunk	chunk
System information (12)	System information (2)	System information (2)		
Chunk free list (1)	Chunk free list (1)	Chunk free list (1)		
Tblspace tblspace (250)	Tblspace tblspace (50)			
Physical log	Free space	Free space		
Logical log				
Sysmaster database				
syutils database				
Free space				

- System information page
 - reserved page
 - 0 page : zero page ()
 - 1 page : config (onconfig)
 - 2,3 page : checkpoint
 - 4,5 page : dbspace
 - 6,7 page : primary chunk
 - 8,9 page : mirror chunk
 - 10,11 page : archive
- 12 Reserved page 가 , 가
- chunk free list page chunk 가
- Tblspace tblspace dbspace tblspace tblspace(table) extnt 가
 - 0 tblspace tblspace bitmap , 1
 - tblspace tblspace , 2



```
Informix> oncheck -pe
DBspace Usage Report: rootdbs Owner: informix Created: 11/16/2001
Chunk Pathname Size Used Free
1 /CHUNK/root.1 10000 9601 399
Description Offset Size
-----
RESERVED PAGES 0 12
CHUNK FREELIST PAGE 12 1
rootdbs: 'informix'.TBLSpace 13 250
PHYSICAL LOG 263 1000
LOGICAL LOG: Log file 1 1263 250
LOGICAL LOG: Log file 2 1513 250
LOGICAL LOG: Log file 3 1763 250
LOGICAL LOG: Log file 4 2013 250
LOGICAL LOG: Log file 5 2263 250
LOGICAL LOG: Log file 6 2513 250
LOGICAL LOG: Log file 7 2763 250
sysmaster: 'informix'.sysdatabases 3013 4
sysmaster: 'informix'.systables 3017 8
```

Space management : creating a space


dbspace / blob space / smart large blob space

- onspaces -c -d dbspace [-t]
 - p chunk -o offset -s
 - [-m chunk offset]
 - onspaces -c -b blobspace -g
 - p chunk -o offset -s
 - [-m chunk offset]
 - onspaces -c -S sbspace
 - p chunk -o offset -s
 - [-m chunk offset]
 - [-Mo offset] [-Ms]
- Kilobyte
- dbspace chunk
 - onstat -d


- Dbspace chunk owner, group informix ,
 - 660 (rw-rw----
 - touch rootdbs
 - chmod 660 rootdbs
 - chown informix:informix rootdbs
 - t temp dbspace , temp
 - dbspace config DBSPACETEMP
 - Space mirror chunk , primary
 - chunk mirror chunk . , mirror
 - chunk config MIRROR가 1
 - Temp dbspace logging
 - Blob space -g blob page ,
 - . Blob space blob
 - switch .
- onmode -l
- Sbspace chunk 가 , -Ms, -Mo 가 가



Sbspace : CLOB/BLOB



Blob space : TEXT/BYTE



```
Informix> onspaces -c -S sbspace -p /CHUNK/sbsp.1 -s 1000 -o 0
Informix> onspaces -c -d tempdbs -t -p /CHUNK/temp.1 -s 1000 -o 0
Informix> onspaces -c -b blobdbs -b 1 /CHUNK/blob.1 -s 1000 -o 0
Informix> onspaces -c -d data1 -p /CHUNK/data1.1 -s 1000 -o 0 -m /CHUNK/mdata1.1 0
Informix> onstat -d
```

Dbspaces							
address	number	flags	fchunk	nchunks	flags	owner	name
a30b7d0	1	0x1	1	1	N	informix	rootdbs
a34a488	2	0x8001	2	1	N S	informix	sbspace
a98fc68	3	0x2001	3	1	N T	informix	tempdbs
a9a02a0	4	0x11	4	1	N B	informix	blobdbs
b937c10	5	0x2	5	1	M	informix	data1

Chunks							
address	chk/dbs	offset	size	free	bpages	flags	pathname
a30b918	1 1	0	10000	537		PO-	/CHUNK/root.1
a34a320	2 2	0	500	336	347	POS	/CHUNK/sbsp.1
		Metadata	100	74	100		
a98fdb0	3 3	0	500	447		PO-	/CHUNK/temp.1
a9a03e8	4 4	0	500	497	500	POB	/CHUNK/blob.1
b937d58	5 5	0	500	447		PO-	/CHUNK/data1.1
b7d42a0	5 5	0	500	0		MO-	/CHUNK/mdata1.1

Space management : adding a chunk / adding, dropping a mirror chunk

- dbspace / blob space / smart large blob space chunk 가
 - onspaces -a space
 - p chunk -o offset -s chunk
 - [-m chunk offset]
 - { [-Mo offset] [-Ms] | -U }
- mirror chunk 가
 - onspaces -m space
 - p chunk -o offset -s chunk
 - m chunk offset [-y]
- mirror chunk
 - onspaces -r space [-y]

```
Informix> onspaces -c -d data2 -p /CHUNK/data2.1 -s 1000 -o 0
Informix> onspaces -m data2 -p /CHUNK/data2.1 -o 0 -m /CHUNK/mdata2.2 0 -y
Informix> onspaces -m rootdbs -p /CHUNK/root.1 -o 0 -m /CHUNK/mroot.1 0 \
-p /CHUNK/root.2 -o 0 -m /CHUNK/mroot.2 0 -y
Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b10b7d0 1 0xa 1 2 MX informix rootdbs
b937ec0 6 0x2 9 1 M informix data2
Chunks
address chk/dbs offset size free bpages flags pathname
b10b918 1 1 0 10000 395 PO- /CHUNK/root.1
b10ba80 1 1 0 10000 0 MX- /CHUNK/mroot.1
b7d4408 6 1 0 500 497 PO- /CHUNK/root.2
b95f570 6 1 0 500 0 MX- /CHUNK/mroot.2
b95f018 9 6 0 500 447 PO- /CHUNK/data2.1
b95f408 9 6 0 500 0 MO- /CHUNK/mdata2.1

Informix> ontape -s
Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b10b7d0 1 0x2 1 2 M informix rootdbs
```

- Space chunk가
 - , space가 mirroring 가 space chunk mirroring
- Sbspace chunk 가 ,
 - U

```
Informix> onspaces -a rootdbs -p /CHUNK/root.2 -s 1000 -o 0
Informix> onspaces -a data1 -p /CHUNK/data1.2 -s 1000 -o 0 -m /CHUNK/mdata1.2 0
Informix> onspaces -a sbspace -p /CHUNK/sbsp.2 -s 1000 -o 0 -U
Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b10b7d0 1 0x1 1 2 N informix rootdbs
b14a758 2 0x8001 2 2 N S informix sbspace
:
b937c10 5 0x2 5 2 M informix data1
Chunks
address chk/dbs offset size free bpages flags pathname
a30b918 1 1 0 10000 537 PO- /CHUNK/root.1
a34a320 2 2 0 500 336 347 POS /CHUNK/sbsp.1
:
Metadata 100 74 100
:
b937d58 5 5 0 500 447 PO- /CHUNK/data1.1
b7d42a0 5 5 0 500 0 MO- /CHUNK/mdata1.1
b7d4408 6 1 0 500 497 PO- /CHUNK/root.2
b7d4570 7 5 0 500 497 PO- /CHUNK/data1.2
b7d46d8 7 5 0 500 0 MO- /CHUNK/mdata1.2
b7d4840 8 2 0 500 497 497 POS /CHUNK/sbsp.2
```

- space mirror 가 , logical log dbspace
 - mirroring , flag가 "M"
 - ontape -s

```
Informix> onspaces -r data1 -y
Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b937c10 5 0x1 5 2 N informix data1
Chunks
address chk/dbs offset size free bpages flags pathname
b937d58 5 5 0 500 447 PO- /CHUNK/data1.1
b7d4570 7 5 0 500 497 PO- /CHUNK/data1.2
```

Space management : changing the chunk state / dropping a space or chunk

```

chunk
  onspaces -s space -p chunk -o offset { -O | -D } [ -y ]

space
  onspaces -d space [ -f ] [ -y ]

chunk
  onspaces -d space -p chunk -o offset [ -f ] [ -y ]
    
```

- I/O 가 chunk 가 chunk가 down ("D") 가 chunk online("O") 가 . (mirror dbspace)

```

Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b937ec0 6 0x2 9 1 M informix data2
Chunks
address chk/dbs offset size free bpages flags pathname
b95f018 9 6 0 500 0 PD- /CHUNK/data2.1
b95f408 9 6 0 500 447 MO- /CHUNK/mdata2.1
Informix> onspaces -s data2 -p /CHUNK/data2.1 -o 0 -0 -y
Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b937ec0 6 0x2 9 1 M informix data2
Chunks
address chk/dbs offset size free bpages flags pathname
b95f018 9 6 0 500 0 PO- /CHUNK/data2.1
b95f408 9 6 0 500 447 MO- /CHUNK/mdata2.1
    
```

- Rootdbs log dbspace(critical dbspace) Space chunk space , chunk mirror chunk가 . Space

```

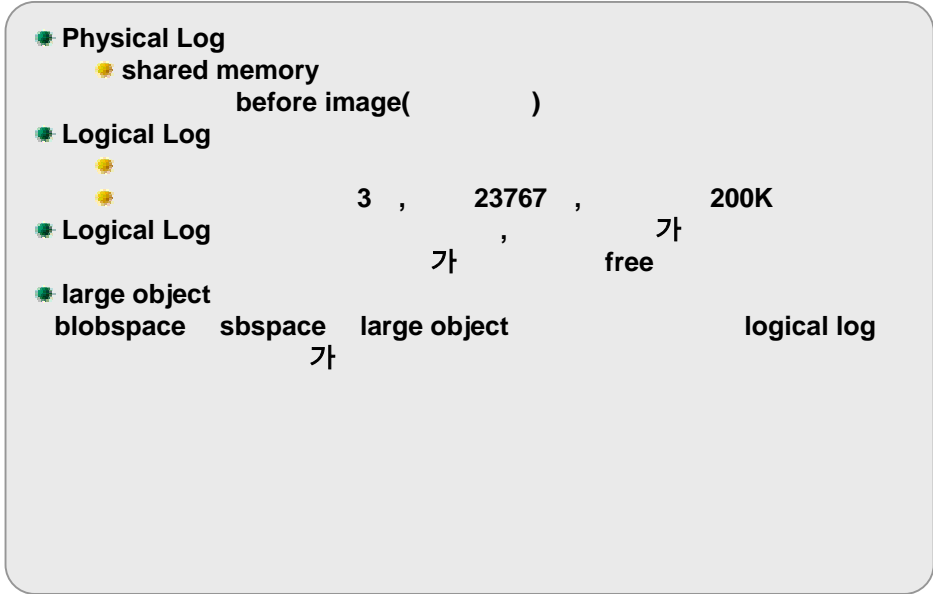
Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b937ec0 6 0x2 9 1 M informix data2
Chunks
address chk/dbs offset size free bpages flags pathname
b95f018 9 6 0 500 0 PD- /CHUNK/data2.1
b95f408 9 6 0 500 447 MO- /CHUNK/mdata2.1
Informix> onspaces -d data2 -y
    
```

```

Informix> onstat -d
Dbspaces
address number flags fchunk nchunks flags owner name
b937c10 5 0x1 5 2 N informix data1
Chunks
address chk/dbs offset size free bpages flags pathname
b937d58 5 5 0 500 447 PO- /CHUNK/data1.1
b7d4570 7 5 0 500 497 PO- /CHUNK/data1.2
Informix> onspaces -d data1 -p /CHUNK/data1.2 -o 0 -y
    
```

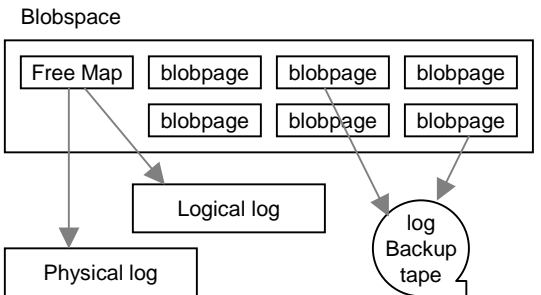





Physical Log / Logical Log



- Physical log rootdbs, dbospace, logical log, physical log, I/O
- dbospace, logical log, DDL, checkpoint, logical log, rootdbs
- Config (\$INFORMIXDIR/etc/\$ONCONFIG) LTAPEDEV /dev/null free
- Blobspace BLOBspace free map blob tracking, blob shared memory buffer, BLOBspace free map, physical log, logical log, blob
- Sbspace smart large object, user가, smart large object 가

4



Log Management

- onstat -l
- logical log 가
 - onparams -a -d dbspace [-s]
- logical log
 - onparams -d -l ID
- onmode -l
- physical log
 - onparams -p -d dbspace -s [-y]
- ontape -s { -N | -B | -U | -A }
 - N : no-logging
 - B : buffered logging
 - U : un-buffered logging
 - A : ANSI mode logging

```
Informix> onstat -l
Physical Logging
Buffer bufused bufsize numpages numwrits pages/io
P-2 0 16 444 42 10.57
phybegin physize phypos phyused %used
100107 1000 33 0 0.00

Logical Logging
Buffer bufused bufsize numrecs numpages numwrits recs/pages pages/io
L-3 0 16 25026 1459 839 17.2 1.7
Subsystem numrecs Log Space used
OLDRSAM 25024 1644548
SBL0B 2 72

address number flags unqid begin size used %used
a07c680 1 U---C-L 1100 1004ef 250 2 0.80
a07c69c 2 U-B---- 1094 1005e9 250 250 100.00
a07c6b8 3 U-B---- 1095 1006e3 250 250 100.00
a07c6d4 4 U-B---- 1096 1007dd 250 250 100.00
a07c6f0 5 U-B---- 1097 1008d7 250 250 100.00
a07c70c 6 U-B---- 1098 1009d1 250 68 27.20
a07c728 7 U-B---- 1099 100acb 250 9 3.60
```

U : Used
 B : backup
 C : currently receiving transactions
 L : contains last completed checkpoint
 A : newly added
 F : Free, available for use

9.2 가 logical log 가, online
 가 quiescent .

```
Informix> onparams -a -d rootdbs -s 200
Informix> onstat -l
address number flags unqid begin size used %used
a07c680 1 U---C-L 1100 1004ef 250 3 1.20
a07c69c 2 U-B---- 1094 1005e9 250 250 100.00
a07c6b8 3 U-B---- 1095 1006e3 250 250 100.00
a07c6d4 4 U-B---- 1096 1007dd 250 250 100.00
a07c6f0 5 U-B---- 1097 1008d7 250 250 100.00
a07c70c 6 U-B---- 1098 1009d1 250 68 27.20
a07c728 7 U-B---- 1099 100acb 250 9 3.60
a07c744 8 A----- 0 60000b 100 0 0.00
Informix > onstape -s
Informix > onstat -l
a07c744 8 F----- 0 60000b 100 0 0.00
```

9.2 “F” logical log 가 “A” 가 , “A”
 “F” 가 , 9.3 가 “A”
 “F” 가 가

Logical log ID unqid log number

```
Informix> onparams -d -l 7 -y
Informix> onparams -d -l 8 -y
Informix> onmode -l
Informix> onparams -p -d data1 -s 200 -y
```

Logical Log 100% switch

physical log 가 가

가

```
Informix> ontape -s -U stores
Informix> ontape -s -B test
```

```
dbaccess sysmaster << !
select name, is_logging, is_buff_log, is_ansi
from sysdatabases ;
!
name stores name test
is_logging 1 is_logging 1
is_buff_log 0 is_buff_log 1
is_ansi 0 is_ansi 0
```



5. PERFORMANCE OPTIMIZATION

Temporary Objects
Parallel Sort / Parallel Index build
Query Optimum Performance
Parallel Database Query
PDQ Administration
Memory Grant Manager
onstat -g mgm
SQL Statement Cache

Temporary Objects

- (Temporary Files)
 - ORDER BY / GROUP BY
 - UNIQUE / DISTINCT
 - Sort merge join
 - index build
- (Temporary tables)
 - Implicit (Select ~ into temp ~)
 - Explicit (Create temp table ~)
 - Blob values / global variables
- 1. PSORT_DBTEMP ()
- 2. DBSPACETEMP
- 3. DBSPACETEMP
- 4-1. /tmp ()
- 4-2. Rootdbs 가 dbspace ()
- Logging WITH NO LOG

- Index order by / group by
- Unique / distinct
- Index sort-merge join
- Hash join hash table
- Index key
- Warm restore logical log
- Into temp
- Create temp table
- Blob stored procedure (global variable)
- temp dbspace
- PDQ round
- robin
- Select ~ into temp rootdbs , create temp table~ 가 dbspace
- Logging DBSPACETEMP , logging DBSPACETEMP logging SQL with no log
- Select ~ into temp ~ with no log;
- Create temp table ~ with no log;

5

```
Informix> onstat -d
```

Dbspaces							
address	number	flags	fchunk	nchunks	flags	owner	name
b10b7d0	1	0x1	1	1	N	informix	rootdbs
b14a758	2	0x8001	2	1	N S	informix	sbspace
b14a8a0	3	0x2001	3	1	N T	informix	tempdbs
b7b8018	4	0x1	4	1	N	informix	data1

Chunks							
address	chk/dbs	offset	size	free	bpages	flags	pathname
b10b918	1	1	0	10000	5154		
b14a320	2	2	0	500	336		
			Metadata	100	74		
b14a488	3	3	0	500	447		
b7b8160	4	4	0	2000	1122		

```
Select * from orders into temp tmp1:
b10b918 1 1 0 10000 5146 PO- /CHUNK/root.1
Select * from orders into temp tmp2 with no log: (DBSPACETEMP not set)
b10b918 1 1 0 10000 5138 PO- /CHUNK/root.1
PO- /CHUNK/data1.1
Select * from orders into temp tmp1 with no log: (DBSPACETEMP=tempdbs)
b14a488 3 3 0 500 439 PO- /CHUNK/temp.1
Create temp table tmp3 (col1 int) with no log: (DBSPACETEMP=tempdbs)
b14a488 3 3 0 500 431 PO- /CHUNK/temp.1
Create temp table tmp3 (col1 int): (DBSPACETEMP not set)
b7b8160 4 4 0 2000 1114 PO- /CHUNK/data1.1
```

Parallel Sort / Parallel Index Build

(parallel sort)

- ORDER BY
- GROUP BY

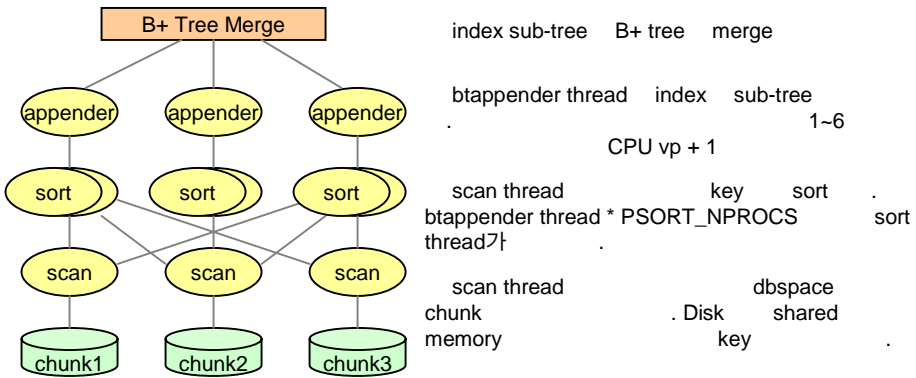
PSORT_NPROCS : sort thread
 PSORT_DBTEMP : sort file
 parallel scan, parallel sort, parallel

index build

- Index Fill Factor
 - FILLFACTOR 80 (\$ONCONFIG)
 - create index ~ fillfactor 80;

PSORT_NPROCS=3

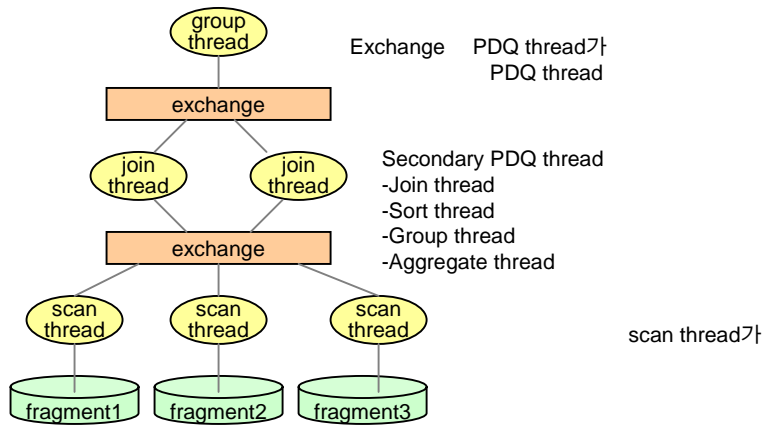
- Sort thread pool, virtual portion sort
- PSORT_NPROCS 0, PSORT_NPROCS 1, PDQ, 2 sort thread가, PDQ가, scan thread가
- PSORT_DBTEMP, DBSPACETEMP, dbspace sort file /tmp
- Index, PDQ, sub-tree 1000, B+ tree (merge), sampling
- Btappender, PSORT_NPROCS, sort 가



Query Optimum Performance

- Query 2가 type
 - Decision Support Queries (DSS)
 - OLTP Queries
- DSS
 - PDQ
 - multi-CPU (fragment)
- PDQ query
- Fragment scan 가

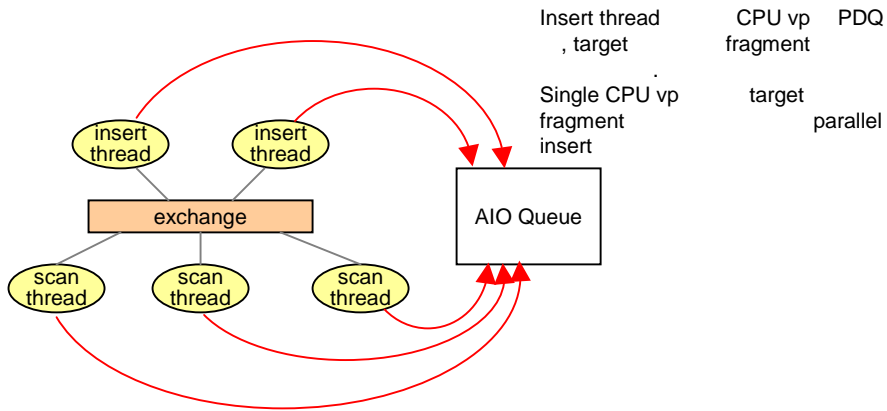
- Decision Support Queries
 - (sequentially)
 - (complex) SQL
- OLTP Queries
 -
 -
 -
- PDQ
 - PDQ (fragment) chunk scan thread가
 - PDQ , CPU 가



Parallel Database Query (PDQ)

- PDQ Operation
 - Parallel scan
 - Parallel join
 - Parallel sort
 - Parallel group
 - Parallel aggregate
- PDQ insert, update, delete
- Parallel insert
 - insert into ~ select ~
 - select ~ into temp ~

- PDQ operation
 - PDQ operation
 - Cursor stability isolation mode
 - For update with hold
 - Correlated subquery parent query
 - SPL query
- Select PDQ가 . DML fetch , fetch
 - insert, update, delete PDQ가 . delete cascade constraint 가 delete PDQ가 .
- Parallel Insert insert thread가
 - Select ~ from
 - Target (Insert) referential constraint 가 (primary key/foreign key)
 - Target 가
 - Target remote
 - Target BLOB
 - Target filtering mode constraint 가
- Select ~ into temp dbspace , with no log DBSPACETEMP logging (light append)
 - buffer pool .



PDQ Administration

(\$ONCONFIG)

- ☀ **MAX_PDQPRIORITY**
- ☀ **DS_MAX_QUERIES**
- ☀ **DS_TOTAL_MEMORY** # Kbytes
- ☀ **DS_MAX_SCANS**
- ☀ **DATASKIP**
- ☀ **OPTCOMPIND**

☀ : export PDQPRIORITY=high

☀ **SQL** : set PDQPRIORITY 50;

☀ **onmode -D** (%) : **MAX_PDQPRIORITY**

☀ **onmode -Q** : **DS_MAX_QUERIES**

☀ **onmode -M** (Kbytes) : **DS_TOTAL_MEMORY**

☀ **onmode -S** : **DS_MAX_SCANS**

• PDQPRIORITY PDQ

PDQPRIORITY		PDQ
OFF	0	PDQ
LOW	1	Parallel Scan only
2 ~ 99		PDQ
HIGH	100	PDQ

- **MAX_PDQPRIORITY** : PDQ query가 PDQ (%)
- **DS_MAX_QUERIES** : PDQ query
- **DS_TOTAL_MEMORY** : PDQ
- **DS_MAX_SCAN** : PDQ scan thread
- **DATASKIP** : off-line chunk
- **OPTCOMPIND** : Optimizer Hint
 - 0 :
 - 1 : repeatable read , 2
 - 2 : query cost가 optimize path
- Onmode PDQ \$ONCONFIG 가 \$ONCONFIG
- PDQ onstat -g mgm

```
Informix> onstat -g mgm
Memory Grant Manager (MGM)
-----
MAX_PDQPRIORITY: 100
DS_MAX_QUERIES: 312
DS_MAX_SCANS: 1048576
DS_TOTAL_MEMORY: 40000 KB
```

Memory Grant Manager (MGM)

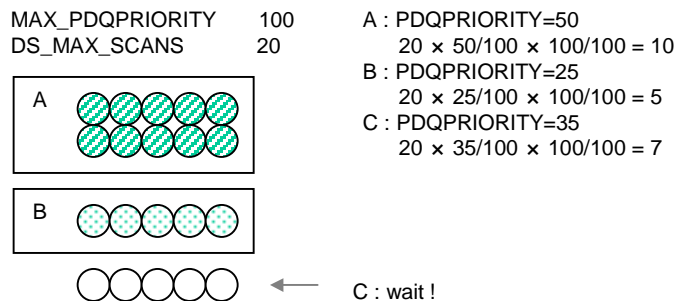

MGM

- scan thread
- quantum : query
- query
- scan thread

PDQ

- PDQ query
- quantum = $DS_TOTAL_MEMORY / DS_MAX_QUERIES$
- total query memory = $DS_TOTAL_MEMORY \times PDQPRIORITY / 100 \times MAX_PDQPRIORITY / 100$
- number of scan threads = $\text{minimum (number of fragments, } DS_MAX_SCANS \times PDQPRIORITY / 100 \times MAX_PDQPRIORITY / 100)$

- PDQPRIORITY 0, memory grant manager가 PDQ
- Total query memory quantum
- MAX_PDQPRIORITY PDQPRIORITY가 100 query가 DS_TOTAL_MEMORY query ready queue



onstat -g mgm

- PDQ query
- PDQ
- PDQ scan thread
- queue
- active ready
- 가

- Active/Ready
- 가

```
Active Queries:
-----
Session  Query      Priority  Thread  Memory  Scans  Gate
   21  aa55260      100    a6fafe0  0/128   0/1    -
Ready Queries:
-----
Session  Query      Priority  Thread  Memory  Scans  Gate
   23  aa12030      100    a7df098  0/128   0/1    1
Free Resource      Average # Minimum #
-----
Memory             113.8 +- 42.7      0
Scans             1048575.0 +- 0.0  1048575
Queries
-----
Average # Maximum # Total #
-----
Active             1.0 +- 0.0      1      9
Ready             1.0 +- 0.0      1      1
Resource/Lock Cycle Prevention count: 0
```

(8K)

-

```
MAX_PDQPRIORITY: 100
DS_MAX_QUERIES: 312
DS_MAX_SCANS: 1048576
DS_TOTAL_MEMORY: 40000 KB
```

- PDQ query

```
Queries:  Active  Ready  Maximum
-         1      1      8
-         Active / Ready
-         가      PDQ query
```

- PDQ

```
Memory:  Total  Free  Quantum
(KB)    40000  0     128
- total : DS_TOTAL_MEMORY
- quantum : DS_TOTAL_MEMORY / DS_MAX_QUERIES
```

- PDQ scan thread

```
Scans:  Total  Free  Quantum
        1048576  1048575  1
- total : DS_MAX_SCANS
-         scan thread scan total - scan free
```

- PDQ queue

```
Load Control: (Memory) (Scans) (Priority) (Max Queries)(Reinit)
               Gate 1 Gate 2 Gate 3 Gate 4 Gate 5
(Queue Length) 1      0      0      0      0
- DS_TOTAL_MEMORY, DS_MAX_SCANS, DS_MAX_QUERIES
```

```
PDQ query
- Priority PDQPRIORITY가 query가 query
- PDQ query가
reinit
query
```



SQL Statement Cache (SSC)

```

($ONCONFIG)
* STMT_CACHE=0 : disabled
* STMT_CACHE=1 : enabled, sessions default off
* STMT_CACHE=2 : enabled, sessions default on
* STMT_CACHE_SIZE : 72KB
( STMT_CACHE가 1 )
* STMT_CACHE=0 : disabled
* STMT_CACHE=1 : enabled
SQL ( STMT_CACHE가 1 )
* set STMT_CACHE off : disabled
* set STMT_CACHE on : enabled
onmode
* onmode -e { ENABLE | ON | OFF | FLUSH }
onstat -g cac stmt
    
```

```

Informix> onstat -g cac stmt
Stmt Cache
Number of lrus : 4
POOLSIZE : 442368
Stmt Cache Entries:
lru hash ref_cnt hits dropped? heap_ptr statement
-----
0 45 0 0 0 a5a1020 stores7:informix.select
* from customer
0 1 0 0 0 a5bf820 stores7:informix.select
* from state
0 120 0 0 0 a58a820 stores7:informix.select
* from call_type
    
```

- cache(SSC) SQL parsing, SQL statement SQL
- SSC DML(select, update, delete, insert) virtual portion
- SSC pool STMT_CACHE_SIZE, SSC cache SQL cache SSC 가
- STMT_CACHE_SIZE 가 SSC
- STMT_CACHE 1 2 enable SSC
- onmode flush
- - EABLE : SSC
- - ON : SSC . Onmode SSC
- - OFF : SSC 가
- - FLUSH : SQL , release flush

6. SERVER MONITORING

Monitoring Utilities

Tables in the sysmaster Database

onstat Options

onstat Multi-thread Options

oncheck Options

Monitor the Message Log

Monitor Shared Memory Segments

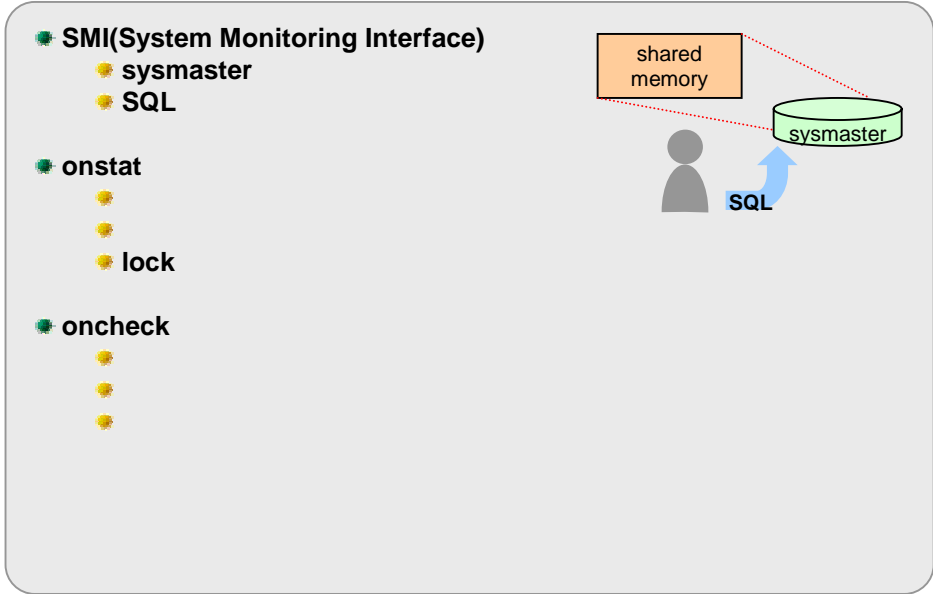
Monitoring Logical Log Space / Chunk Usage

Monitoring User Activity

Identifying Server Resources

Monitoring Lock Usage

Monitoring Utilities



- SMI(System Monitoring Interface)
 - sysmaster pointing virtual SQL read only
 - sysmaster 가
 - sysmaster 가 lock isolation 가
 - insert, update, delete 가
 - dbschema, dbexport 가
 - rowid select

- onstat
 - (interactive) : onstat -i
 - : onstat -g sub_option
 - : onstat --
 - : onstat -
 - : onstat -r []

```
Informix> onstat -i
> p
> r 5 interactive mode Ctrl-c
> u
```

- oncheck
 - 가 oncheck (, oncheck -pT, oncheck -pt) shared lock 가
 - 가
 - -c sub , -p
 - sub

Tables in the sysmaster Database

sysdatabases	
systabnames	
syslogs	logical log . size used logical log full
sysdbspaces	dbspace
syschunks	chunk . nfree
syslocks	lock
sysvpprof	가 (Virtual Processor)
sysessions	
sysesprof	가
sysextents	extent
syschkio	chunk I/O
sysptprof	가 table tblspace . memory tblspace close shared sysptprof 가 free
sysprofile	checkpoint, rollback
sysadtnfo	auditing
sysaudit	16 audit mask. onaudit , , 가
sysconfig	
sysdri	(data replication)
syssewts	

```

1)
SELECT name FROM sysdatabases

2)
SELECT tablename FROM systabnames WHERE dbsname = "db_name"

3)
SELECT uniqid, ROUND ((used/size) * 100, 2) FROM syslogs;

4) chunk (onstat -d)
SELECT chknum, (100 * (chksize - nfree)/chksize) FROM syschunks

5) online chunk
SELECT c.chknum chunk_number, d.name dbspace, 'On-Line' status
FROM syschunks c, sysdbspaces d
WHERE c.dbsnum = d.dbsnum AND c.is_offline = 0;

6) lock process id
SELECT ss.username || '-' || ss.sid user_session, sp.locksheld, ss.pid
FROM sysesprof sp, sysessions ss WHERE sp.sid = ss.sid;

7) VP
SELECT * FROM sysmaster.sysessions ss, sysmaster.sysesprof sp
WHERE sp.sid = dbinfo ("sessionid") AND sp.sid = ss.sid;

8) lock (onstat -u, onstat -k)
SELECT username, sid, waiter, dbsname, tablename, rowidlk, keynum, type
FROM syslocks, sysessions WHERE sysessions.sid = syslocks.owner

9) extent
SELECT tablename, count(*) num_extents FROM sysextents
WHERE dbsname = 'db_name' GROUP BY tablename ORDER BY tablename;

```

onstat Options

usage: onstat [-abcdghklmpstuzBDFRX][-i] [-r seconds] [-o file] [infile]

--	onstat
-a	onstat -cuskbtdlp
-b	, resident
-B	
-c	\$ONCONFIG
-C	B-tree cleaner
-d	dbspace chunk
-D	chunk page
-f	DATASKIP dbspace
-F	flush
-h	
-i	onstat interactive mode
-k	lock
-K	byte-range lock
-l	physical log logical log
-m	20
-o filename	filename shared memory
-O	Optical subsystem memory cache
-p	profile
-P	partition
-r seconds	. 5
-R	LRU queue
-s	latch
-t	tblspace
-u	
-x	
-X	
-z	onstat profile 0 reset

onstat Multi-thread Options

usage: onstat -g options

ath	
wai	waiting
act	active
rea	ready
sle	sleeping
spi	spin lock
sch	VP
lmx	lock mutex
wmx	mutex waiter
con	condition waiter
stk	<tid>
glo	Multi-thread global information
mem	<pool name session id>
seg	
rbm	resident block map
nbm	non-resident block map
afr	<pool name session id> pool
ffr	<pool name session id> free pool
ufr	<pool name sid> pool
iovp	VP
iof	chunk
ioq	queue
iob	IO VP big buffer
ppf	[<partition number> 0] active partition profile
tpf	[<tid> 0] thread profile
ntu	network profile
ntt	network

ntm	network
ntd	thread
nss	<session id> shared memory connection
nsc	<client id> shared memory connection
nsd	shared memory
sts	
dic	dictionary cache
qst	queue
wst	wait
prc	
dsc	
ses	<session id>
sql	<session id> SQL
dri	

oncheck Options

Usage: oncheck [-clist] [-plist] [-qny] [{ database:[owner.]table | TBLSpace number | Chunk number } { rowid | page number }]

cc	system catalog
cd	simple large object
cD	-cd blob
ce	chunk free list free space, tblspace extent, smart large object extent metadata
ci	index key
cl	-ci index rowid
cr	root dbspace reserved page
cR	-cr physical log, logical log ,
cs	sbspace smart large object meta data
n	index 가
pB	blobpage
pc	system catalog
pd	16
pD	blob 16
pe	chunk tblspace extent
pk	index key
pK	index key rowid
pl	index leaf node key
pL	index leaf node key rowid
pp	
pP	-pp chunk rowid
pr	-cr
pR	-cR
ps	sbspace smart large object
pS	-ps , sbspace extent

pt	fragment tblspace
pT	-pt index
q	checking validation
x	index lock shared
y	index 가

Monitor the Message Log

Message Log (\$ONCONFIG MSGPATH)

- Fast recovery
- Checkpoint
- shared memory
- I/O

- log 가 message
- \$ONCONFIG MSGPATH 가 . message log 가
- onstat message log 20
 - cat /dev/null > online.log
 - onstat -m

```


12:30:58 Checkpoint Completed: duration was 1 seconds
12:30:58 Continuing Long Transaction (for COMMIT):tx 0x80b3d8
username:informix uid:200
12:32:52 Logical Log 20 Complete
12:34:53 Checkpoint Completed: duration was 0 seconds
12:35:27 dynamically allocated new shared memory segment (size 8388608)
12:37:07 Checkpoint Completed: duration was 3 seconds
12:40:36 Checkpoint Completed: duration was 0 seconds
12:41:19 asf_shm_be.c,line 1949,thread 23, failure, system out of semaphores.
12:41:19 PANIC: Attempting to bring system down
12:42:07 INFORMIX-OnLine Initialized -- Shared Memory Initialized
12:42:07 Physical Recovery Started
12:42:07 Physical Recovery Complete: 0 Pages Restored
12:42:07 Logical Recovery Started
12:42:07 Logical Recovery allocating 3 worker threads
('OFF_RECVRY_THREADS').
12:42:12 Logical Recovery Complete
7 Committed,2 Rolled Back, 0 Open, 0 Bad Locks 12:42:12 On-Line Mode

```

Monitor Shared Memory Segments

- Shared memory segment
 - onstat -g seg (logical)
 - ipcs (physical)
- Shared memory segment (\$ONCONFIG)
 - SHMVIRTSIZE
 - SHMADD
 - SHMTOTAL
- Shared memory segment
 - onmode -a <seg_size>
- Shared memory
 - onstat -g mem
- shared memory segment
 - onstat -F

- Shared memory segment
- 가 shared memory segment 가 , kernel shared memory
- SHMVIRTSIZE shared memory virtual portion 가 SHMADD , onmode -a virtual portion shared memory segment가 SHMTOTAL



SHMVIRTSIZE 20000
 SHMADD 10000
 SHMTOTAL 95000 onmode -a 20000

```
Informix> onstat -g seg
Segment Summary:
(resident segments are not locked)
id key addr size ovhd class blkused blkfree
32 1381451777 800000 987136 372 R 117 4
33 1381451778 8f1000 4096000 252 V 333 167
34 1381451779 cd9000 1048576 204 M 106 22
```

8K

```
Informix> onstat -g mem
Pool Summary:
name class addr totalsize freesize #allocfrag #freefrag
resident R c001a010 13516800 10104 2 1
global V c0d0a010 237568 19384 425 46
mt V c0d0e010 1941504 134264 912 40
aio V c0d2e010 950272 117016 207 36
dictpool V c0d32010 65536 7824 48 4
procpool V c0d36010 16384 7080 12 2
shmcon M e74c2010 229376 16248 2 1
btclean V c0f5a010 8192 3400 8 1
aio_fpf V C0f7C010 81920 16248 2 1
13 V b350020 65536 20504 92 16
15 V b367020 57344 7056 110 10
```

```
Informix> onstat -g ses
session #RSAM total used
id user tty pid hostname threads memory
15 informix 23 16072 kor-oliv 1 57344 50288
13 informix 21 15798 kor-oliv 1 65536 45032
```


Monitoring Logical Log Space / Chunk Usage

Logical Log

- onstat -l
- SMI (syslogs, systrans)

Chunk Usage

- onstat -d
- SMI (syschunks)

- logical log free 가 . 가
- Logical log가 (free log)가 가 hang 가
- logical log가 message log 가
- SMI logical log

```
select uniqid,(used/size*100) from syslogs
where uniqid >= (select min(tx_loguniq) from systrans where tx_loguniq > 0)
union
select uniqid,0.00 from syslogs
where uniqid < (select min(tx_loguniq) from systrans where tx_loguniq > 0);
```

- chunk
 - dbspace chunk가 full chunk 가
 - dbspace 가
 - dbspace chunk가 off-line 가
- SMI chunk

```
SELECT chknum, dbsnum FROM syschunks WHERE is_offline = 1 or mis_offline = 1
```

```
SELECT chknum, dbsnum FROM syschunks WHERE nfree < 100
```

```
Informix> onstat -l
Physical Logging
Buffer bufused bufsize numpages numwrits pages/io
P-2 0 16 444 42 10.57
phybegin physize phypos phyused %used
100107 1000 33 0 0.00
Logical Logging
Buffer bufused bufsize numrecs numpages num
L-3 0 16 25026 1459 839
Subsystem numrecs Log Space used
OLDRSAM 25024 1644548
SBL0B 2 72
```

address	number	flags	uniqid	begin	size	used	%used
a07c680	1	U---C-L	1100	1004ef	250	2	0.80
a07c69c	2	U-B----	1094	1005e9	250	250	100.00
a07c6b8	3	U-B----	1095	1006e3	250	250	100.00
a07c6d4	4	U-B----	1096	1007dd	250	250	100.00
a07c6f0	5	U-B----	1097	1008d7	250	250	100.00
a07c70c	6	U-B----	1098	1009d1	250	68	27.20
a07c728	7	U-B----	1099	100acb	250	9	3.60

B = Backed up
 C = Currently receiving transactions
 U = In use
 A = Newly added
 L = Contains last checkpoint



Monitoring User Activity - 1

- onstat -g ses
- onstat -g ses <session_id>

SQL

- onstat -g sql

User Activity

- Session
- lock
- SQL
-
-
-
-
-
- onstat -g ses pid가 0

```
Informix> onstat -g ses 19
```

session	id	user	tty	pid	hostname	#RSAM threads	total memory	used memory
19		informix	-	17230	kor-oliv	1	65536	47536

tid	name	rstcb	flags	curstk	status
39	sqlexec	ad18818	Y-BP---	1648	cond wait(netnorm)

name	class	addr	totalsize	freesize	#allocfrag	#freefrag
19	V	b31f020	65536	18000	110	17

name	free	used	name	memory pool
overhead	0	1648	scb	
opentable	0	5200	filetable	0 928
ru	0	224	log	0 2152
temprec	0	10056	keys	0 352
ralloc	0	6144	gentcb	0 1248
ostcb	0	2520	sqscb	0 12216
sql	0	40	rdahead	0 112
hashfiletab	0	280	osenv	0 1464
sqtcdb	0	2640	fragman	0 216

Sess Id	SQL Stmt type	Current Database	Iso Lvl	Lock Mode	SQL ERR	ISAM ERR	F.E. Vers
19	-	stores_demo	CR	Not Wait	0	0	9.03

Last parsed SQL statement :

```
select customer.*, orders.order_num from customer, outer orders where customer.customer_num = orders.customer_num into temp custs_orders
```


User-created Temp tables :

partnum	tabname	rowsize
100088	custs_orders	138
100087	cust_orders	138

```
Informix> onstat -g ses
```

session id	user	tty	pid	hostname	#RSAM threads	total memory	used memory
20	informix	-	0	-	0	12288	7656
19	informix	14	17230	kor-oliv	1	65536	47536
3	informix	-	0	-	0	1	shared memory
2	informix	-	0	-	0	1	shared memory

```
Informix> onstat -g sql
```

Sess Id	SQL Stmt type	Current Database	Iso Lvl	Lock Mode	SQL ERR	ISAM ERR	F.E. Vers
21	SELECT	stores_demo	CR	Not Wait	0	0	9.03
19	-	stores_demo	CR	Not Wait	0	0	9.03

Monitoring User Activity - 2

- **user thread lock**
 - onstat -u
 - onstat -k
- **SMI user thread**
 - select * from sysseprof
 - select * from sysssessions
- **select longtxs, logspused, maxlogsp from sysseprof where sid = 34**
- **onstat -g tpf <session_id>**
- **onmode -z <session_id>**

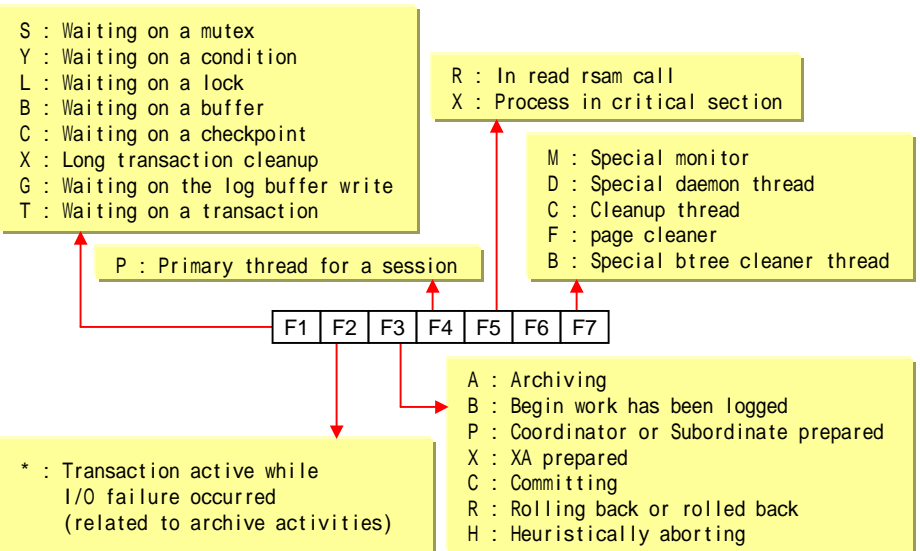
• onstat -u flag
onstat -g ses

```
Informix> onstat -u
```

address	flags	sessid	user	tty	wait	tout	locks	nreads	nwrites
a120018	---P--D	1	informix	-	0	0	0	50	33
a1204bc	---P--F	0	informix	-	0	0	0	0	278
a120e04	---P--B	7	informix	-	0	0	0	1	0
a1212a8	---P--D	9	informix	-	0	0	0	0	0
a12174c	L--PR--	26	joeg	ttyp2	a01f8d4	-1	1	0	0
a121bf0	Y-BP---	23	lizg	ttyp1	a2fb050	0	6	0	0

```
Informix> onstat -k
```

address	wtlist	owner	lklist	type	tblsnum	rowid	key#/bsiz
a01f838	0	a121bf0	0 HDR+S	100002	205	0	
a01f86c	0	a121bf0	a01f838 HDR+IX	100085	0	0	
a01f8d4	a12174c	a121bf0	a01fa40 HDR+X	100085	400	0	
a01f908	0	a121bf0	a01f970 HDR+X	100085	200	2	
a01f970	0	a121bf0	a01f86c HDR+X	100085	100	1	
a01fa40	0	a121bf0	a01f908 HDR+X	100085	300	3	
a01fbe0	0	a12174c	0 S	100002	205	0	



• logical log
가
- longtxs : long transaction
- logspused : logical log (bytes)
- maxlogsp : logical log (bytes)
가
onmode -z disconnect
resource

Identifying Server Resources

- onstat -p
- Lock (contention)
- Dead lock time-out
- Read cache

```
Informix> onstat -p
dskreads pagreads bufreads %cached dskwrits pagwrits bufwrits %cached
216 260 3305 93.46 49 71 158 68.99
isamtot open start read write rewrite delete commit rollbk
1408 114 194 603 70 0 0 2 1
gp_read gp_write gp_rewrt gp_del gp_alloc gp_free gp_curs
2 0 0 0 0 0 2
ovlock ovuserthead ovbuff usercpu syscpu numckpts flushes
0 0 0 3.13 0.76 6 3404
bufwaits lokwaits lockreqs deadlks dltouts ckpwais compress seqscans
15 0 1300 0 0 0 3 14
ixda-RA idx-RA da-RA RA-pgsused lchwaits
11 0 16 25 0
```

```
ovlock ovuserthead ovbuff usercpu syscpu numckpts flushes
0 0 0 3.13 0.76 6 3404
```

over lock LOCKS 가
over buffer가 BUFFERS 가

```
bufwaits lokwaits lockreqs deadlks dltouts ckpwais compress seqscans
15 0 1300 0 0 0 3 14
```

(lockwaits/lockreqs) * 100 should be < 1%

```
bufwaits lokwaits lockreqs deadlks dltouts ckpwais compress seqscans
15 0 1300 0 0 0 3 14
```

dltouts DEADLOCK_TIMEOUT 가

```
dskreads pagreads bufreads %cached dskwrits pagwrits bufwrits %cached
216 260 3305 93.46 49 71 158 68.99
```

```
ixda-RA idx-RA da-RA RA-pgsused lchwaits
```

```
11 0 16 25 0
```

RA-pgsused ixda-RA + idx-RA + da-RA 가

- ovlock, ovuserthead, ovbuff 가 0가
- Lock request ,
 - page level locking row level locking
 - locking
 - Repeatable read isolation level
 - dirty read isolation 가
- dead lock time-out remote
 - , dltouts 가 . dead lock time-out
 - remote 가 .(60)
- DEADLOCK_TIMEOUT 가 .(60)
- Local deal lock deadlks 가 , rollback
 - . deadlks 가 .
- Read cache 가 BUFFERS 가 . Read
 - cache 95%
- Read Ahead가 read cache , bufwaits

Monitoring Lock Usage

- onstat -k
- Lock type
 - shared lock / exclusive lock / update lock / intent lock
- Lock
 - database / table / page / row / byte / key
- Lock
 - Logging : transaction (commit rollback lock)
 - Non-logging : SQL
- Dynamic Lock Manager
 - LOCKS lock 가
 - 100,000 lock table 16 가
- Smart Large Object byte range locking
 - smart large object locking 가
 - lock smart large object locking
 - BLOB lock mode : onspaces -DfLOCK_MODE=BLOB (default)
 - Range lock mode : onspaces -DfLOCK_MODE=RANGE

- Shared lock read concurrency 가 shared lock 가
- Exclusive lock update concurrency 가
- Update lock update cursor 가 update concurrency exclusive lock
- Intent lock shared lock update exclusive lock promotion 가 shared lock
- Repeatabe read lock , key delete insert shared exclusive

```
Informix> onstat -k
Locks
address wtlst owner lklist type tbslnum rowid key#/bsiz
a05408c 0 ad18818 0 HDR+S 100002 207 0
a05408c 0 ad18818 0 HDR+S 100270 0 0
a S : shared lock
a U : update lock
a X : exclusive lock
a IS : intent shared lock
a IX : intent exclusive lock
a SIX : Shared intent exclusive lock
a SR : shared repeatable read lock
a XR : exclusive repeatable read lock
a054134 0 ad18818 a0540e0 HDR+SR 100088 203 K- 1
a054188 0 ad18818 a054134 HDR+SR 100088 205 K- 1
```

Annotations in the image:

- database lock (points to rowid 207)
- table lock (points to rowid 100)
- row lock (points to rowid 200)
- page lock (points to rowid 805)
- key lock (points to key# 1)

```
lock
- tbslnum 100002 (tblspace=0x00100002) lock
lock
SMI rowid
select name from sysdatabases
where hex(rowid) = '0x00000207'
lock
rowid 0 , tbslnum 100002가
SMI tbslnum
select tablename from systabnames
where hex(partnum) = '0x00100273'
- rowid slot 00
lock 00가 tbslnum 100002가
row lock
- Lock mode가 row key key lock
, key#/bsiz K 가 . (Lock mode가 page
page lock )
- varchar byte lock type
"B"
```

7. *BACKUP & RESTORE*

Backup

Logical Log Backup

Restore

ontape Utility

ontape Example

onbar Utility

Restore

- physical restore
 - logical restore
- cold restore
 - warm restore
 - mixed restore
- Log Salvage

- Physical restore 0, 1, 2 dbspace, blob space, logical restore, physical restore, logical log
- dbspace rootdbs(critical db space) 가
 - cold restore off-line
 - cold restore critical db space 가
 - warm restore on-line
 - 가
 - cold restore critical db space (rootdbs, log db space) (mixed restore)
 - warm restore
 - 가 off-line restore logical log
 - log salvage cold restore
 - on-line log
 - 가 logical restore

ontape Example

- **ontape -s**
- **ontape -a**
- **Warm restore**
 - **ontape -r -D dbs2**
- **Cold restore**
 - **ontape -r**

```
Informix> ontape -s
Please enter the level of archive to be performed (0, 1, or 2) 0
Please mount tape 1 on /dev/tape and press Return to continue ...
10 percent done.
20 percent done.
30 percent done.
40 percent done.
100 percent done.
Please label this tape as number 1 in the arc tape sequence.
This tape contains the following logical logs:
 66
Program over.
```

`ontape -s -L 0`

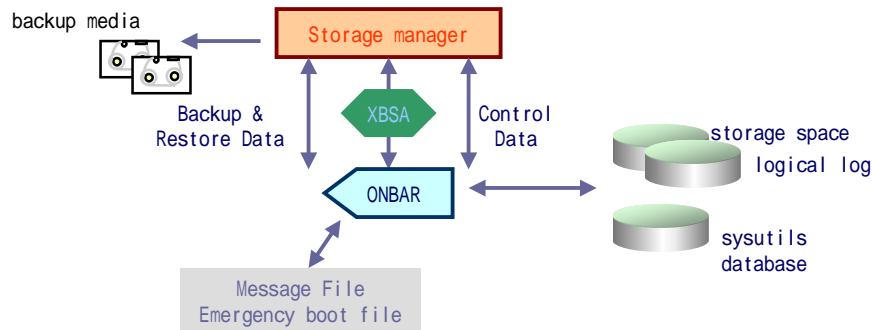
```
Informix> ontape -a
Performing automatic backup of logical logs.
Please mount tape 1 on /dev/ltape and press Return to continue ...
Do you want to back up the current logical log? (y/n) y
Please label this tape as number 1 in the log tape sequence.
This tape contains the following logical logs:
 66
Program over.
```

```
Informix> ontape -r -D dbs2
Please mount tape 1 on /tmp/tape and press Return to continue ...
Archive Tape Information
Tape type:      Archive Backup Tape
Online version: Informix Dynamic Server Version 9.30.UC1G1
Archive date:   Mon Jul 15 10:49:05 2002
User id:       informix
Terminal id:    /dev/pts/21
Archive level: 0
Tape device:   /dev/tape
Tape blocksize (in k): 16
Tape size (in k): 1024000
Tape number in series: 1
Continue restore? (y/n) y
Spaces to restore: 1 [dbs2]
Restore a level 1 archive (y/n) n
Do you want to restore log tapes? (y/n) y
Roll forward should start with log number 66
Please mount tape 1 on /dev/ltape and press Return to continue ...
Do you want to restore another log tape? (y/n) n
Program over.
```

```
Informix> ontape -r
Please mount tape 1 on /tmp/tape and press Return to continue ...
:
Archive Information
Informix Dynamic Server Copyright(C) 1986-1999 Informix Software, Inc.
Initialization Time      12/14/2001 11:42:46
System Page Size        2048
Version                  12
Archive CheckPoint Time 07/15/2002 10:49:05
Dbspaces
number  flags  fchunk  nchunks  flags  owner      name
1       20001  1       1        N      informix   rootdbs
2       8001   2       1        N S    informix   sbospace
3       20001  3       1        N      informix   datadbs
Chunks
chk/dbs offset  size  free  bpages  flags  pathname
1  1  0    25000  1063  PO-    /DBS/rootdbs_chk1
2  2  0    5000   267   POS    /DBS/sbospace_chk1
3  3  0    5000   1409  PO-    /DBS/datadbs_chk1
Continue restore? (y/n) y
Do you want to back up the logs? (y/n) n
Restore a level 1 archive (y/n) n
Do you want to restore log tapes? (y/n) n
Program over.
```

onbar Utility

- automatic(on-demand) continuous
- Third party storage manager XBSA
- `onbar -b [-L level] [-f list_file | dbspace_list | -w]`
- `onbar -b -l [-c | -s | -C]`
- `onbar -r [dbspace_list | -f list_file | -w | -t time | -n log]`



- onbar X/Open Backup Services Application Programming Interface (XBSA) storage manager communicate RDBMS 가 .
- onbar ontape 가
가
- operation
- storage manager 가
- dbspace
- 가 (whole system backup)
- sysutils onbar가 object instance
가 catalog
- bar_server
- bar_object
- bar_action
- bar_instance
- `onbar -b` : dbspace, blobspace
- `onbar -b dbspace1 dbspace2` : dbspace1, dbspace2
- `onbar -b -L 1 : 1`
- `onbar -b -L 2 dbspace1 dbspace2 : 2` : 2 dbspace1, dbspace2
- `onbar -b -f file_name` : file_name dbspace, blobspace
- `onbar -b -w` : whole system backup ()
- `onbar -b -w -L 1 : 1` : whole system backup
- `onbar -b -l` :
- `onbar -b -l -c` :
- \$ONCONFIG ALARMPROGRAM log_full.sh log_full.sh
onbar -b -l 가 가 .
- `onbar -b -l -s` : cold restore log salvage 가
- `onbar -r` : dbspace, blobspace ()
- `onbar -r -w` : whole system backup ()
- `onbar -r -f dbspace1 dbspace2` : dbspace1, dbspace2
- `onbar -r -t 2002-05-30 14:20:00` :
- `onbar -r -n 64` : 64