

Operating System

Solaris X86-84 v10.10

Virtual processors: 16

Memory: 16GB (locked in VMware)

Paging space: 8GB (after increase from 512MB)

VMware Multipathing: disabled

XtremIO disks: RAW device mappings (RDM)

/etc/system non-default settings:

set zfs:zfs_vdev_max_pending = 32

set zfs:zfs_nocacheflush = 1

set zfs:zfs_arc_max = 104857600

set zfs:zfs_arc_min = 104857600

set project.max-shm-memory = 4294967295

Oracle

Database: 11.2.0.4.0 with ASM

Non-default init.ora settings (based on XtremIO best practice):

*._db_block_prefetch_limit=0

*._db_block_prefetch_quota=0

*._db_file_noncontig_mblock_read_count=0

*._disk_sector_size_override=TRUE

*.db_file_multiblock_read_count=16

*.pga_aggregate_target=512M

*.processes=1000

*.sga_max_size=3G

*.sga_target=3G

SLOB config (initial settings):

Slob version 2 with patch update 2014.03.11. Wait kit disabled (outcommented in runit.sh).

UPDATE_PCT=0

RUN_TIME=300

WORK_LOOP=0

SCALE=900000

WORK_UNIT=256

REDO_STRESS=LIGHT

LOAD_PARALLEL_DEGREE=2

SHARED_DATA_MODULUS=0

Total SLOB blocks: 900000 * 64 users = 57,600,000 blocks of 8KB each = 450000 MB

Disk configuration

ASM (before loading SLOB data):

SYS:xtremdb> @report-asmdisks

Disk Group Name Path File Name Fail Group File Size (MB) Used Size (MB) Pct. Used

DATA /dev/rdsk/c3t0d0s6 DATA_0000 DATA_0000 133,088 600 .45

```

/dev/rdsk/c3t1d0s6 DATA_0001 DATA_0001 133,088 592 .44
/dev/rdsk/c3t2d0s6 DATA_0002 DATA_0002 133,088 584 .44
/dev/rdsk/c3t3d0s6 DATA_0003 DATA_0003 133,088 592 .44
/dev/rdsk/c3t4d0s6 DATA_0004 DATA_0004 133,088 616 .46
*****
665,440 2,984
REDO1 /dev/rdsk/c3t5d0s6 REDO1_0000 REDO1_0000 12,256 272 2.22
/dev/rdsk/c3t6d0s6 REDO1_0001 REDO1_0001 12,256 272 2.22
*****
24,512 544
REDO2 /dev/rdsk/c3t8d0s6 REDO2_0000 REDO2_0000 12,256 80 .65
/dev/rdsk/c3t9d0s6 REDO2_0001 REDO2_0001 12,256 56 .46
*****
24,512 136
-----
Grand Total: 714,464 3,664

```

Note that REDO2 was not used.

Creating tablespace for SLOB:

```

CREATE BIGFILE TABLESPACE "IOPS" DATAFILE '+DATA(DATAFILE)' SIZE 450G REUSE
AUTOEXTEND ON NEXT 100M MAXSIZE UNLIMITED NOLOGGING EXTENT
MANAGEMENT LOCAL SEGMENT SPACE MANAGEMENT AUTO
/

```

Redo logs (with 4K block size emulation enabled, XtremIO best practice):

```

SYS:xtremdb > @ls-redologs
GRP THREAD SEQ ARC STATUS BLKSZ SIZE_MB FILENAME
-----

```

```

4 1 246 NO CURRENT 4096 128 +REDO1/logfile01.rdo
5 1 244 NO INACTIVE 4096 128 +REDO1/logfile02.rdo
6 1 245 NO INACTIVE 4096 128 +REDO1/logfile03.rdo

```

```

SYS:xtremdb > @ls-tablespaces

```

```

NAME TOTAL MB FREE MB STATUS CONTENTS AUT INC MAX_MB
-----

```

```

IOPS 465920 5035 ONLINE PERMANENT YES 12800 512000
SYSAUX 540 32 ONLINE PERMANENT YES 1280 32767.9844
SYSTEM 750 9 ONLINE PERMANENT YES 1280 32767.9844
USERS 5 4 ONLINE PERMANENT YES 160 32767.9844
TEMP ONLINE TEMPORARY
UNDOTBS1 285 257 ONLINE UNDO YES 640 32767.9844

```

6 rows selected.

ZFS configuration

```

bash-3.2# zpool status

```

```

pool: data_pool

```

```

state: ONLINE

```

```

scan: none requested

```

```

config:

```

```
NAME STATE READ WRITE CKSUM
data_pool ONLINE 0 0 0
c2t0d0 ONLINE 0 0 0
c2t1d0 ONLINE 0 0 0
c2t2d0 ONLINE 0 0 0
c2t3d0 ONLINE 0 0 0
c2t4d0 ONLINE 0 0 0
errors: No known data errors
pool: online_redo1
state: ONLINE
scan: none requested
config:
```

```
NAME STATE READ WRITE CKSUM
online_redo1 ONLINE 0 0 0
c2t5d0 ONLINE 0 0 0
logs
c2t6d0 ONLINE 0 0 0
errors: No known data errors
pool: online_redo2
state: ONLINE
scan: none requested
config:
```

```
NAME STATE READ WRITE CKSUM
online_redo2 ONLINE 0 0 0
c2t8d0 ONLINE 0 0 0
logs
c2t9d0 ONLINE 0 0 0
errors: No known data errors
bash-3.2# zpool list
```

```
NAME SIZE ALLOC FREE CAP HEALTH ALTROOT
data_pool 595G 466G 129G 78% ONLINE -
online_redo1 9.94G 174M 9.77G 1% ONLINE -
online_redo2 9.94G 408K 9.94G 0% ONLINE -
bash-3.2# zpool get all data_pool
```

(listed only non-default and relevant settings)

```
NAME PROPERTY VALUE SOURCE
data_pool size 595G -
data_pool capacity 78% -
data_pool free 129G -
data_pool allocated 466G -
data_pool readonly off -
bash-3.2# zfs get all data_pool/data
```

(listed only non-default and relevant settings)

```
NAME PROPERTY VALUE SOURCE
data_pool/data type filesystem -
data_pool/data creation Fri Aug 1 13:13 2014 -
data_pool/data used 462G -
data_pool/data available 120G -
data_pool/data referenced 462G -
```

data_pool/data compressratio 1.00x -
data_pool/data recordsize 8K local
data_pool/data mountpoint /data_pool/data default
data_pool/data checksum on default
data_pool/data compression off local
data_pool/data atime off inherited from data_pool
data_pool/data primarycache all local
data_pool/data secondarycache all default
data_pool/data logbias throughput inherited from data_pool
data_pool/data sync standard default

After moving IOPS tablespace to ZFS, list of datafiles:

SYS:xtremdb > @ls-datafiles

FILE# STATUS MB BS NAME

1 SYSTEM 750 8192 +DATA/xtremdb/datafile/system.256.854450867
2 ONLINE 540 8192 +DATA/xtremdb/datafile/sysaux.257.854450867
3 ONLINE 285 8192 +DATA/xtremdb/datafile/undotbs1.258.854450867
4 ONLINE 5 8192 +DATA/xtremdb/datafile/users.259.854450867
5 ONLINE 465920 8192 /data_pool/data/iops.dbf

FILE# STATUS MB BS NAME

1 ONLINE 1024 8192 +DATA/xtremdb/tempfile/temp.268.854450967

GRP# STATUS MB BS NAME

4 CURRENT 128 4096 +REDO1/logfile01.rdo
5 INACTIVE 128 4096 +REDO1/logfile02.rdo
6 INACTIVE 128 4096 +REDO1/logfile03.rdo