ASM HANDS-ON TRAINING

Lab 6 Looking at The Partnership Status Metadata

Alejandro Vargas | Principal Support Consultant
Oracle Advanced Customer Services

INDEX

What is The Partnership and Status Table (PST)

From Note 344927.1

“PST - Partnership and Status Table contains the status information about the ASM disks in a disk group - disk number, status (either online or offline), partner disk number, failure group info (11g) and heartbeat info. AU number 1 in every disk within a disk group is reserved for PST. Only a few disk actually have a PST - in external redundancy group we only have one PST table, in normal redundancy group (double mirroring) we have up to 3 PST and in high redundancy we have up to 5 PST. The GMON process is responsible for PST processing. See kfdp source code for more informaiton.

In the first block on AU "1" is the PST header.“

Looking on PST Metadata

We have a Normal Redundancy Diskgroup set; we will lookup at the PST Metadata on their four ASM disks

The PST must to be replicated on 2*f+1 to tolerate a failure of at most f disks.

Our NR diskgroup was built using 4 disks: VOL1, 2,3,4

Reading the PST Using Kfed

PST ON VOL1

```
[oracle@asmxpt ~]$ kfed read /dev/oracleasm/disks/VOL1
aunum=1
kfbh.endian:                          1 ; 0x000: 0x01
kfbh.hard:                          130 ; 0x001: 0x82
kfbh.type:                           17 ; 0x002:
KFBTYP_PST_META
```

PST ON VOL3

```
[oracle@asmxpt ~]$ kfed read /dev/oracleasm/disks/VOL3
aunum=1
kfbh.endian:                          1 ; 0x000: 0x01
kfbh.hard:                          130 ; 0x001: 0x82
kfbh.type:                           17 ; 0x002:
KFBTYP_PST_META
```

```
kfbh.datfmt:                          1 ; 0x003: 0x01
kfbh.block.blk:                     256 ; 0x004: T=0
NUMB=0x100
kfbh.block.obj:            2147483648 ; 0x008: TYPE=0x8
NUMB=0x0
kfbh.check:                4105465801 ; 0x00c: 0xf4b46fc9
kfbh.fcn.base:                        0 ; 0x010: 0x00000000
kfbh.fcn.wrap:                        0 ; 0x014: 0x00000000
kfbh.spare1:                          0 ; 0x018: 0x00000000
kfbh.spare2:                          0 ; 0x01c: 0x00000000
kfdpHdrB.time.hi:              32917704 ; 0x000: HOUR=0x8
DAYS=0x6 MNTH=0x2 YEAR=0x7d9
kfdpHdrB.time.lo:          1951769600 ; 0x004: USEC=0x0
MSEC=0x169 SECS=0x5 MINS=0x1d
kfdpHdrB.last:                        1 ; 0x008: 0x00000001
kfdpHdrB.next:                        1 ; 0x00c: 0x00000001
kfdpHdrB.copyCnt:                     2 ; 0x010: 0x02
kfdpHdrB.ub1spare:                    0 ; 0x011: 0x00
kfdpHdrB.ub2spare:                    0 ; 0x012: 0x0000
kfdpHdrB.incarn:                      0 ; 0x014: 0x00000000
kfdpHdrB.copy[0]:                     0 ; 0x018: 0x0000
kfdpHdrB.copy[1]:                     2 ; 0x01a: 0x0002
kfdpHdrB.copy[2]:                     0 ; 0x01c: 0x0000
kfdpHdrB.copy[3]:                     0 ; 0x01e: 0x0000
kfdpHdrB.copy[4]:                     0 ; 0x020: 0x0000
kfdpHdrB.dtaSz:                       4 ; 0x022: 0x0004
ub1[0]:                               2 ; 0x024: 0x02
ub1[1]:                               0 ; 0x025: 0x00
```

There is no PST on VOL2

```
[oracle@asmxpt ~]$ kfed read /dev/oracleasm/disks/VOL2
aunum=1
kfbh.endian:                          1 ; 0x000: 0x01
kfbh.hard:                          130 ; 0x001: 0x82
kfbh.type:                           13 ; 0x002:
KFBTYP_PST_NONE
kfbh.datfmt:                          1 ; 0x003: 0x01
kfbh.block.blk:            2147483648 ; 0x004: T=1
NUMB=0x0
kfbh.block.obj:            2147483649 ; 0x008: TYPE=0x8
NUMB=0x1
kfbh.check:                  17662464 ; 0x00c: 0x010d8200
kfbh.fcn.base:                        0 ; 0x010: 0x00000000
```

```
kfbh.datfmt:                          1 ; 0x003: 0x01
kfbh.block.blk:                     256 ; 0x004: T=0
NUMB=0x100
kfbh.block.obj:            2147483650 ; 0x008: TYPE=0x8
NUMB=0x2
kfbh.check:                4105465803 ; 0x00c: 0xf4b46fcb
kfbh.fcn.base:                        0 ; 0x010: 0x00000000
kfbh.fcn.wrap:                        0 ; 0x014: 0x00000000
kfbh.spare1:                          0 ; 0x018: 0x00000000
kfbh.spare2:                          0 ; 0x01c: 0x00000000
kfdpHdrB.time.hi:              32917704 ; 0x000: HOUR=0x8
DAYS=0x6 MNTH=0x2 YEAR=0x7d9
kfdpHdrB.time.lo:          1951769600 ; 0x004: USEC=0x0
MSEC=0x169 SECS=0x5 MINS=0x1d
kfdpHdrB.last:                        1 ; 0x008: 0x00000001
kfdpHdrB.next:                        1 ; 0x00c: 0x00000001
kfdpHdrB.copyCnt:                     2 ; 0x010: 0x02
kfdpHdrB.ub1spare:                    0 ; 0x011: 0x00
kfdpHdrB.ub2spare:                    0 ; 0x012: 0x0000
kfdpHdrB.incarn:                      0 ; 0x014: 0x00000000
kfdpHdrB.copy[0]:                     0 ; 0x018: 0x0000
kfdpHdrB.copy[1]:                     2 ; 0x01a: 0x0002
kfdpHdrB.copy[2]:                     0 ; 0x01c: 0x0000
kfdpHdrB.copy[3]:                     0 ; 0x01e: 0x0000
kfdpHdrB.copy[4]:                     0 ; 0x020: 0x0000
kfdpHdrB.dtaSz:                       4 ; 0x022: 0x0004
ub1[0]:                               2 ; 0x024: 0x02
ub1[1]:                               0 ; 0x025: 0x00
```

There is no PST on VOL4

```
[oracle@asmxpt ~]$ kfed read /dev/oracleasm/disks/VOL4
aunum=1
kfbh.endian:                          1 ; 0x000: 0x01
kfbh.hard:                          130 ; 0x001: 0x82
kfbh.type:                           13 ; 0x002:
KFBTYP_PST_NONE
kfbh.datfmt:                          1 ; 0x003: 0x01
kfbh.block.blk:            2147483648 ; 0x004: T=1
NUMB=0x0
kfbh.block.obj:            2147483651 ; 0x008: TYPE=0x8
NUMB=0x3
kfbh.check:                  17662466 ; 0x00c: 0x010d8202
kfbh.fcn.base:                        0 ; 0x010: 0x00000000
```

```
kfbh.fcn.wrap:                    0 ; 0x014: 0x00000000
kfbh.spare1:                      0 ; 0x018: 0x00000000
kfbh.spare2:                      0 ; 0x01c: 0x00000000
```

```
kfbh.fcn.wrap:                    0 ; 0x014: 0x00000000
kfbh.spare1:                      0 ; 0x018: 0x00000000
kfbh.spare2:                      0 ; 0x01c: 0x00000000
```

PST Breakdown

kfbh.endian
kf3.h /* endianness of writer */
Little endian = 1
Big endian = 0

```
kfbh.endian:                      1 ; 0x000: 0x01
```

kfbh.hard
kf3.h /* H.A.R.D. magic # and block size */

```
kfbh.hard:                      130 ; 0x001: 0x82
```

kfbh.type
kf3.h /* metadata block type */

```
kfbh.type:                       17 ; 0x002: KFBTYP_PST_META
```

kfbh.datfmt
kf3.h /* metadata block data format */

```
kfbh.datfmt:                      1 ; 0x003: 0x01
```

kfbh.block
kf3.h /* block location of this block */
blk -- Disk header should have T=0 and NUMB=0x0
obj -- Disk header should have TYPE=0x8 NUMB=<disknumber>

blk and obj values are derived from a series of macros in kf3.h. See "KFBL Macros" in kf3.h for more information.

```
kfbh.block.blk:                       256 ; 0x004: T=0 NUMB=0x100
kfbh.block.obj:                2147483648 ; 0x008: TYPE=0x8 NUMB=0x0
```

kfbh.check
kf3.h /* check value to verify consistency */

```
kfbh.check:                    4105465801 ; 0x00c: 0xf4b46fc9
```

kfbh.fcn
kf3.h /* change number of last change */

```
kfbh.fcn.base:                          0 ; 0x010: 0x00000000
kfbh.fcn.wrap:                          0 ; 0x014: 0x00000000
```

kfdpHdrB.time.hi
kf3.h Hi ordered bits from the last committed PST update

```
kfdpHdrB.time.hi:               32917704 ; 0x000: HOUR=0x8 DAYS=0x6 MNTH=0x2 YEAR=0x7d9
```

kfdpHdrB.time.lo
kf3.h Low ordered bits from the last committed PST update

```
kfdpHdrB.time.lo:             1951769600 ; 0x004: USEC=0x0 MSEC=0x169 SECS=0x5 MINS=0x1d
```

kfdpHdrB.last
kf3.h /* last version number */

```
kfdpHdrB.last:                          1 ; 0x008: 0x00000001
```

kfdpHdrB.next
kf3.h /* next version number */

```
kfdpHdrB.next:                          1 ; 0x00c: 0x00000001
```

kfdpHdrB.copyCnt
kf3.h /* # of PST copies */
This defaults to "1" for external redundancy, "3" for normal redundancy and "5" for high redundancy. If the number of failure groups is less than the default value, the number failure groups is the value used.

```
kfdpHdrB.copyCnt:                       2 ; 0x010: 0x02
```

kfdpHdrB.incarn
kf3.h /* incarnation of <copy> */
This is set to kfdpHdrB.last when the PST is moved to another disk.

```
kfdpHdrB.incarn:                        0 ; 0x014: 0x00000000
```

kfdpHdrB.copy[0-4]
kf3.h /* disks holding the PST copies */
[0] -- external redundancy
[0-2] -- normal redundancy
[0-4] -- high redundancy

```
kfdpHdrB.copy[0]:                       0 ; 0x018: 0x0000
kfdpHdrB.copy[1]:                       2 ; 0x01a: 0x0002
kfdpHdrB.copy[2]:                       0 ; 0x01c: 0x0000
kfdpHdrB.copy[3]:                       0 ; 0x01e: 0x0000
kfdpHdrB.copy[4]:                       0 ; 0x020: 0x0000
```

kfdpHdrB.dtaSz
kf3.h /* # dta entries in PST */
This is the number of disks that it needs to keep track of.

```
kfdpHdrB.dtaSz:                          4 ; 0x022: 0x0004
```

ub1[0-4027]

End of Lab6