

- Linux Kernel Parameter - Dirty ; Page cache control** 3
- Overview** 3
- Definition** 3
- Process** 3
- Opinion** 4
- Latest Issue and Kernel Parameters** 4
- Details of Mysql_Cache_Unmap** 4
- Attention** 6

Linux Kernel Parameter - Dirty ; Page cache control

: http://kakaodbe.blogspot.kr/2014/03/linux-kernel-parameters-and-mysql_2226.html

Overview

MySQL

Definition

Linux /proc/sys/vm , Page Cache , dirty data , disk , flush , 6

Kernel Parameter	description	Default
dirty_background_ratio	dirty page , pdflush	10(%)
dirty_ratio	dirty page , process가 dirty page	40(%)
dirty_background_bytes	dirty_background_ratio (dirty_background_ratio) , 0 , 0 , dirty_background_ratio 가	0
dirty_bytes	dirty_ratio (dirty_ratio) , 0 , dirty_ratio 가	0
dirty_writeback_centisecs	page flush dirty	500(1/100sec)
dirty_expire_centisecs	page , dirty page가	3000(1/100sec)

Process

pdflush 'dirty_writeback_centisecs' writeback ,
 page cache dirty page
 'dirty_background_ratio' 'dirty_expire_centisecs'
 dirty page disk flush .

Opinion

dirty_background_ratio

I/O . pdfflush disk flush dirty page ,
 I/O I/O
 가 dirty_background_ratio
 . (DB insert update !)

dirty_writeback_centisecs

page 가 page

dirty_expire_centisecs

30 , disk
 30 dirty data가 disk
 . (InnoDB innodb_flush_method
 innodb_max_dirty_pages_pct
 .) ** innodb_max_dirty_pages_pct : innodb buffer fool dirty page

Latest Issue and Kernel Parameters

i/o 가 가 가
 Filesystem Cache dirty data , swap
 cache unmap 가 dirty page 1 가 Cache
 Cache data dirty data disk i/o가 가

```
/proc/sys/vm/dirty_background_ratio : 10 ==> 1
/proc/sys/vm/dirty_expire_centisecs : 3000 ==> 1000
```

dirty_background_ratio 1 cache 1% dirty page가
 , dirty_expire_centisecs 10 dirty page가 cache

Details of Mysql_Cache_Unmap

Kakao DB Team Filesystem Cache swap cache_unmap
 . (url :
<http://kakaodbe.blogspot.kr/2013/09/mysql-linux-file-system-cache-2.html>) cache
 unmap

```
db          mysql_cache_unmap          Matt          가
            . (                    : mysql_cache_unmap.c)
cache unmapping          Crontab          .
```

```
# Crontab Setting
*/10 * * * * root LD_LIBRARY_PATH=/otp/mysql/lib:
/otp/mysql/admin/mysql_cache_unmap --defaults-file=/etc/my.cnf --
binary_os_cache_size=1024M > /otp/mysql/admin/mysql_cache_unmap.log2>&1
```

```
-defaults-file          mysql server  configuration files          . cache unmap
                        binary log  data file          innodb redo log
                        . -binary_os_cache_size          Linux OS Cache
cache_unmap            (          1024M)  Linux OS
Cache
mysql_cache_unmap
```

```
root@host:~ 14:10:09> cat /opt/mysql/admin/mysql_cache_unmap.log
Read configuration
innodb_data_dir : /opt/mysql/data
innodb_log_dir : /opt/mysql/data
binary_log_dir : /opt/mysql/data/mysql-binary
binary_os_cache_size : 1073741824
relay_log_dir : /opt/mysql/data/mysql-relay
> unmap_file_all : datafile : /opt/mysql/data/dbname1/tablename1.ibd
> unmap_file_all : datafile : /opt/mysql/data/dbname1/tablename2.ibd
> unmap_file_all : logfile : /opt/mysql/data/ib_logfile0
> unmap_file_all : logfile : /opt/mysql/data/ib_logfile1
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041728 : 0 ~
60083119
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041727 : 0 ~
104857974
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041726 : 0 ~
104857975
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041725 : 0 ~
104857883
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041724 : 0 ~
104857852
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041723 : 0 ~
104857920
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041722 : 0 ~
104857989
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041721 : 0 ~
104857781
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041720 : 0 ~
104857784
> skip unmap file: binary : /opt/mysql/data/mysql-binary.041719 : 0 ~
104857666
```



```
#include <fcntl.h>
int unmap_file_segment(const char *fpath, size_t start, size_t len){
    int fd = open(fpath, O_RDONLY);
    if (fd < 0){
        fprintf(stderr, "ERROR : Failed to open %s\n", fpath);
        return 1;
    }
    int r = posix_fadvise(fd, start, len, POSIX_FADV_DONTNEED);
    if (r != 0){
        fprintf(stderr, "ERROR : posix_fadvise failed for %s\n", fpath);
    }
    close(fd);
    /* if posix_fadvise is succeeded, then sleep 25 milli seconds */
    usleep(25 * 1000);
    return 0;
}
```

advice		'POSIX_FADV_DONTNEED'			
advice	가	, 'POSIX_FADV_DONTNEED'			page
cache(filesystem cache)		file	offset	len	
unmapping					

~~DISCUSSION~~

From:
<https://atl.kr/dokuwiki/> - AllThatLinux!

Permanent link:
https://atl.kr/dokuwiki/doku.php/linux_kernel_parameter_-_dirty

Last update: **2016/03/05 23:34**

