

LCOV - code coverage report

Current view: [directory](#) - [fs/ext4](#) - [mballoc.h](#) ([source](#) / [functions](#))

Found Hit C

Test: [kernel_2_6_31_ext4_round_3.info](#)

Lines:

2	2	1
---	---	---

Date: [2009-10-24](#)

Functions:

0	0	
---	---	--

```
1      : /*
2      : *   fs/ext4/mballoc.h
3      : *
4      : *   Written by: Alex Tomas <alex@clusterfs.com>
5      : *
6      : */
7      : #ifndef _EXT4_MBALLOC_H
8      : #define _EXT4_MBALLOC_H
9      :
10     : #include <linux/time.h>
11     : #include <linux/fs.h>
12     : #include <linux/namei.h>
13     : #include <linux/quotaops.h>
14     : #include <linux/buffer_head.h>
15     : #include <linux/module.h>
16     : #include <linux/swap.h>
17     : #include <linux/proc_fs.h>
18     : #include <linux/pagemap.h>
19     : #include <linux/seq_file.h>
20     : #include <linux/version.h>
21     : #include <linux/blkdev.h>
22     : #include <linux/mutex.h>
23     : #include "ext4_jbd2.h"
24     : #include "ext4.h"
25     :
26     : /*
27     : * with AGGRESSIVE_CHECK allocator runs consistency checks over
28     : * structures. these checks slow things down a lot
29     : */
30     : #define AGGRESSIVE_CHECK__
31     :
32     : /*
33     : * with DOUBLE_CHECK defined mballoc creates persistent in-core
34     : * bitmaps, maintains and uses them to check for double allocations
35     : */
36     : #define DOUBLE_CHECK__
37     :
38     : /*
39     : *
40     : * #define MB_DEBUG__
41     : * #ifdef MB_DEBUG
42     : *     #define mb_debug(fmt, a...)    printk(fmt, ##a)
43     : * #else
44     : *     #define mb_debug(fmt, a...)
45     : * #endif
46     : *
47     : */
48     : * with EXT4_MB_HISTORY mballoc stores last N allocations in memory
49     : * and you can monitor it in /proc/fs/ext4/<dev>/mb_history
50     : */
51     : #define EXT4_MB_HISTORY
52     : #define EXT4_MB_HISTORY_ALLOC      1      /* allocation */
```

```

53 : #define EXT4_MB_HISTORY_PREALLOC      2      /* preallocated blocks used */
54 : #define EXT4_MB_HISTORY_DISCARD      4      /* preallocation discarded */
55 : #define EXT4_MB_HISTORY_FREE          8      /* free */
56 :
57 : #define EXT4_MB_HISTORY_DEFAULT        (EXT4_MB_HISTORY_ALLOC | \
58 :                                         EXT4_MB_HISTORY_PREALLOC)
59 :
60 : /*
61 :  * How long mballoc can look for a best extent (in found extents)
62 :  */
63 : #define MB_DEFAULT_MAX_TO_SCAN        200
64 :
65 : /*
66 :  * How long mballoc must look for a best extent
67 :  */
68 : #define MB_DEFAULT_MIN_TO_SCAN        10
69 :
70 : /*
71 :  * How many groups mballoc will scan looking for the best chunk
72 :  */
73 : #define MB_DEFAULT_MAX_GROUPS_TO_SCAN  5
74 :
75 : /*
76 :  * with 'ext4_mb_stats' allocator will collect stats that will be
77 :  * shown at umount. The collecting costs though!
78 :  */
79 : #define MB_DEFAULT_STATS                1
80 :
81 : /*
82 :  * files smaller than MB_DEFAULT_STREAM_THRESHOLD are served
83 :  * by the stream allocator, which purpose is to pack requests
84 :  * as close each to other as possible to produce smooth I/O traffic
85 :  * We use locality group prealloc space for stream request.
86 :  * We can tune the same via /proc/fs/ext4/<partition>/stream_req
87 :  */
88 : #define MB_DEFAULT_STREAM_THRESHOLD    16      /* 64K */
89 :
90 : /*
91 :  * for which requests use 2^N search using buddies
92 :  */
93 : #define MB_DEFAULT_ORDER2_REQS        2
94 :
95 : /*
96 :  * default group prealloc size 512 blocks
97 :  */
98 : #define MB_DEFAULT_GROUP_PREALLOC      512
99 :
100 :
101 : struct ext4_free_data {
102 :     /* this links the free block information from group_info */
103 :     struct rb_node node;
104 :
105 :     /* this links the free block information from ext4_sb_info */
106 :     struct list_head list;
107 :
108 :     /* group which free block extent belongs */
109 :     ext4_group_t group;
110 :
111 :     /* free block extent */
112 :     ext4_grpblk_t start_blk;
113 :     ext4_grpblk_t count;
114 :
115 :     /* transaction which freed this extent */
116 :     tid_t    t_tid;
117 : };

```

```

118 :
119 : struct ext4_prealloc_space {
120 :     struct list_head    pa_inode_list;
121 :     struct list_head    pa_group_list;
122 :     union {
123 :         struct list_head pa_tmp_list;
124 :         struct rcu_head pa_rcu;
125 :     } u;
126 :     spinlock_t          pa_lock;
127 :     atomic_t            pa_count;
128 :     unsigned            pa_deleted;
129 :     ext4_fsblk_t        pa_pstart;    /* phys. block */
130 :     ext4_lblk_t        pa_lstart;    /* log. block */
131 :     unsigned short      pa_len;      /* len of preallocated chunk */
132 :     unsigned short      pa_free;     /* how many blocks are free */
133 :     unsigned short      pa_type;     /* pa type. inode or group */
134 :     spinlock_t          *pa_obj_lock;
135 :     struct inode        *pa_inode;   /* hack, for history only */
136 : };
137 :
138 : enum {
139 :     MB_INODE_PA = 0,
140 :     MB_GROUP_PA = 1
141 : };
142 :
143 : struct ext4_free_extent {
144 :     ext4_lblk_t fe_logical;
145 :     ext4_grpblk_t fe_start;
146 :     ext4_group_t fe_group;
147 :     int fe_len;
148 : };
149 :
150 : /*
151 :  * Locality group:
152 :  *   we try to group all related changes together
153 :  *   so that writeback can flush/allocate them together as well
154 :  *   Size of lg_prealloc_list hash is determined by MB_DEFAULT_GROUP_PREALLOC
155 :  *   (512). We store prealloc space into the hash based on the pa_free blocks
156 :  *   order value, ie, fls(pa_free)-1;
157 :  */
158 : #define PREALLOC_TB_SIZE 10
159 : struct ext4_locality_group {
160 :     /* for allocator */
161 :     /* to serialize allocates */
162 :     struct mutex    lg_mutex;
163 :     /* list of preallocations */
164 :     struct list_head lg_prealloc_list[PREALLOC_TB_SIZE];
165 :     spinlock_t      lg_prealloc_lock;
166 : };
167 :
168 : struct ext4_allocation_context {
169 :     struct inode *ac_inode;
170 :     struct super_block *ac_sb;
171 :
172 :     /* original request */
173 :     struct ext4_free_extent ac_o_ex;
174 :
175 :     /* goal request (after normalization) */
176 :     struct ext4_free_extent ac_g_ex;
177 :
178 :     /* the best found extent */
179 :     struct ext4_free_extent ac_b_ex;
180 :
181 :     /* copy of the best found extent taken before preallocation efforts */
182 :     struct ext4_free_extent ac_f_ex;

```

```

183 :
184 :     /* number of iterations done. we have to track to limit searching */
185 :     unsigned long ac_ex_scanned;
186 :     __u16 ac_groups_scanned;
187 :     __u16 ac_found;
188 :     __u16 ac_tail;
189 :     __u16 ac_buddy;
190 :     __u16 ac_flags;           /* allocation hints */
191 :     __u8 ac_status;
192 :     __u8 ac_criteria;
193 :     __u8 ac_repeats;
194 :     __u8 ac_2order;           /* if request is to allocate 2^N blocks and
195 :                                * N > 0, the field stores N, otherwise 0 */
196 :     __u8 ac_op;               /* operation, for history only */
197 :     struct page *ac_bitmap_page;
198 :     struct page *ac_buddy_page;
199 :     /*
200 :      * pointer to the held semaphore upon successful
201 :      * block allocation
202 :      */
203 :     struct rw_semaphore *alloc_semp;
204 :     struct ext4_prealloc_space *ac_pa;
205 :     struct ext4_locality_group *ac_lg;
206 : };
207 :
208 : #define AC_STATUS_CONTINUE      1
209 : #define AC_STATUS_FOUND        2
210 : #define AC_STATUS_BREAK        3
211 :
212 : struct ext4_mb_history {
213 :     struct ext4_free_extent orig; /* orig allocation */
214 :     struct ext4_free_extent goal; /* goal allocation */
215 :     struct ext4_free_extent result; /* result allocation */
216 :     unsigned pid;
217 :     unsigned ino;
218 :     __u16 found; /* how many extents have been found */
219 :     __u16 groups; /* how many groups have been scanned */
220 :     __u16 tail; /* what tail broke some buddy */
221 :     __u16 buddy; /* buddy the tail ^^^ broke */
222 :     __u16 flags;
223 :     __u8 cr:3; /* which phase the result extent was found at */
224 :     __u8 op:4;
225 :     __u8 merged:1;
226 : };
227 :
228 : struct ext4_buddy {
229 :     struct page *bd_buddy_page;
230 :     void *bd_buddy;
231 :     struct page *bd_bitmap_page;
232 :     void *bd_bitmap;
233 :     struct ext4_group_info *bd_info;
234 :     struct super_block *bd_sb;
235 :     __u16 bd_blkbits;
236 :     ext4_group_t bd_group;
237 :     struct rw_semaphore *alloc_semp;
238 : };
239 : #define EXT4_MB_BITMAP(e4b) ((e4b)->bd_bitmap)
240 : #define EXT4_MB_BUDDY(e4b) ((e4b)->bd_buddy)
241 :
242 : #ifndef EXT4_MB_HISTORY
243 : static inline void ext4_mb_store_history(struct ext4_allocation_context *ac)
244 : {
245 :     return;
246 : }
247 : #endif

```

```
248 :
249 : #define in_range(b, first, len) ((b) >= (first) && (b) <= (first) + (len) - 1)
250 :
251 : static inline ext4_fsblk_t ext4_grp_offs_to_block(struct super_block *sb,
252 :                                                    struct ext4_free_extent *fex)
253 : {
254 :     ext4_fsblk_t block;
255 :
256 690609042 :     block = (ext4_fsblk_t) fex->fe_group * EXT4_BLOCKS_PER_GROUP(sb)
257 :                                     + fex->fe_start
258 :                                     + le32_to_cpu(EXT4_SB(sb)->s_es->s_first_data_block);
259 230203014 :     return block;
260 : }
261 : #endif
```

Generated by: [LCOV version 1.8](#)