

94


```

181         actAG_free = bmap_recptr->free_blocks - inactAG_free;
182         avg_free = actAG_free / num_active_AGs;
183
184         if ((bmap_recptr->bmpctl_bufptr->dn_agpref > highest_active_AG) || (bmap_recptr->bmpctl_bufptr->d
n_agfree[bmap_recptr->bmpctl_bufptr->dn_agpref] < avg_free)) {
185             /* preferred AG is not valid */
186             if (avg_free == 0) {
187                 bmap_recptr->bmpctl_bufptr->dn_agpref = 0;
188             } else {
189                 bmap_recptr->bmpctl_bufptr->dn_agpref = -1;
190                 for (agidx = 0; ((agidx < MAXAG) &&
191                     (bmap_recptr->bmpctl_bufptr->
192                     dn_agpref < 0)); agidx++) {
193                     if (bmap_recptr->AGFree_tbl[agidx] >= avg_free) {
194                         bmap_recptr->bmpctl_bufptr->
195                         dn_agpref = agidx;
196                     }
197                 }
198             }
199         }
200         aglevel = BMAPSZTOLEV(sb_ptr->s_agsize);
201         l2nl =
202             agg_recptr->log2_blkspereag - (L2BPERDMAP +
203             aglevel * L2LPERCTL);
204         agheight = l2nl >> 1;
205         agwidth = 1 << (l2nl - (agheight << 1));
206         for (index = 5 - agheight, agstart = 0, n = 1; index > 0;
207             index--) {
208             agstart += n;
209             n <<= 2;
210         }
211
212         bmap_recptr->bmpctl_bufptr->dn_aglevel = aglevel;
213         bmap_recptr->bmpctl_bufptr->dn_agheight = agheight;
214         bmap_recptr->bmpctl_bufptr->dn_agwidth = agwidth;
215         bmap_recptr->bmpctl_bufptr->dn_agstart = agstart;
216         bmap_recptr->bmpctl_bufptr->dn_agl2size =
217             agg_recptr->log2_blkspereag;
218         bmap_recptr->bmpctl_bufptr->dn_agsize = sb_ptr->s_agsize;
219         bmap_recptr->bmpctl_bufptr->dn_maxfreebud = max_buddy;
220
221         for (agidx = 0; (agidx < MAXAG); agidx++) {
222             /* rebuild the free list */
223             bmap_recptr->bmpctl_bufptr->dn_agfree[agidx] =
224                 bmap_recptr->AGFree_tbl[agidx];
225         }
226
227         /* swap if on big endian machine */
228         ujfs_swap_dbmap(bmap_recptr->bmpctl_bufptr);
229
230         /*
231          * write the updated control page back onto the device
232          */
233         bcr_rc = mapctl_put((void *) bmap_recptr->bmpctl_bufptr);
234     }
235     return (bcr_rc);
236 }
237
238 /*****
239  * NAME: ctlpage_verify
240  *
241  * FUNCTION: Verify the control page of the filesystem block map.
242  *
243  * PARAMETERS:
244  *     max_buddy - input - the data value which should be in the root
245  *                   of the highest Lx page in the map.
246  *
247  * RETURNS:
248  *     success: FSCK_OK
249  *     failure: something else
250  */
251 int ctlpage_verify(int8_t max_buddy)
252 {
253     int bcv_rc = FSCK_OK;
254     unsigned agidx;
255     int32_t max_level = 0;
256     int32_t highest_active_AG = 0;
257     int32_t l2nl, n, index, aglevel, agheight, agwidth, agstart;
258
259     bcv_rc = mapctl_get(bmap_recptr->bmpctl_agg_fsblk_offset,
260         (void **) &(bmap_recptr->bmpctl_bufptr));
261
262     if (bcv_rc == FSCK_OK) {
263         /* got the control page in the buffer */
264
265         /* swap if on big endian machine */
266         ujfs_swap_dbmap((struct dbmap *) bmap_recptr->bmpctl_bufptr);
267
268         if (bmap_recptr->bmpctl_bufptr->dn_mapsize != bmap_recptr->total_blocks) {
269             /* bad number of blocks in the aggregate */

```

```

270         bmap_recptr->ctl_other_error = -1;
271         fsck_send_msg(fsck_BMAPCASB, 0);
272     }
273     if (bmap_recptr->bmapctl_bufptr->dn_nfree != bmap_recptr->free_blocks) {
274         /* bad number of free blocks in the aggregate */
275         bmap_recptr->ctl_other_error = -1;
276         fsck_send_msg(fsck_BMAPCNF, 0);
277     }
278     if (bmap_recptr->bmapctl_bufptr->dn_l2nbperpage != agg_recptr->log2_blksperspg) {
279         /* bad log2( blocks per page ) */
280         bmap_recptr->ctl_other_error = -1;
281         fsck_send_msg(fsck_BMAPCL2BPP, 0);
282     }
283     if (bmap_recptr->bmapctl_bufptr->dn_numag != agg_recptr->num_ag) {
284         /* bad number of alloc groups */
285         bmap_recptr->ctl_other_error = -1;
286         fsck_send_msg(fsck_BMAPCNAG, 0);
287     }
288     max_level = BMAPSZTOLEV(agg_recptr->sb_agg_fsblk_length);
289     if (bmap_recptr->bmapctl_bufptr->dn_maxlevel != max_level) {
290         /* bad maximum block map level */
291         bmap_recptr->ctl_other_error = -1;
292         fsck_send_msg(fsck_BMAPCMXLVL, 0);
293     }
294     for (agidx = 0; (agidx < MAXAG); agidx++) {
295         /* check out the active AGs */
296         if (bmap_recptr->AGActive[agidx]) {
297             highest_active_AG = agidx;
298         }
299     }
300     /*
301     * format does not include blocks allocated to the bad block inode
302     * when it determines the dn_maxag. Subsequent activity may or
303     * may not have altered dn_maxag based on blocks allocated to the
304     * bad block inode. All we know for sure is that dn_maxag shouldn't
305     * be larger than appropriate for the highest allocated block.
306     */
307     if (bmap_recptr->bmapctl_bufptr->dn_maxag > highest_active_AG) {
308         /* bad highest active alloc group */
309         bmap_recptr->ctl_other_error = -1;
310         fsck_send_msg(fsck_BMAPCMAAG, 0);
311     }
312     if (bmap_recptr->bmapctl_bufptr->dn_aggpref > highest_active_AG) {
313         /* bad preferred alloc group */
314         bmap_recptr->ctl_other_error = -1;
315         fsck_send_msg(fsck_BMAPCPAG, 0);
316     }
317     aglevel = BMAPSZTOLEV(sb_ptr->s_agsize);
318     l2nl =
319         agg_recptr->log2_blkspersag - (L2BPERDMAP +
320         aglevel * L2LPERCTL);
321     agheight = l2nl >> 1;
322     agwidth = 1 << (l2nl - (agheight << 1));
323     for (index = 5 - agheight, agstart = 0, n = 1; index > 0;
324         index--) {
325         agstart += n;
326         n <= 2;
327     } /* end for index */
328
329     if (bmap_recptr->bmapctl_bufptr->dn_aglevel != aglevel) {
330         /* bad level holding an AG */
331         bmap_recptr->ctl_other_error = -1;
332         fsck_send_msg(fsck_BMAPCDMCLAG, 0);
333     }
334     if (bmap_recptr->bmapctl_bufptr->dn_agheight != agheight) {
335         /* bad dmapctl height holding an AG */
336         bmap_recptr->ctl_other_error = -1;
337         fsck_send_msg(fsck_BMAPCDMCLH, 0);
338     }
339     if (bmap_recptr->bmapctl_bufptr->dn_agwidth != agwidth) {
340         /* bad width at level holding an AG */
341         bmap_recptr->ctl_other_error = -1;
342         fsck_send_msg(fsck_BMAPCDMCLW, 0);
343     }
344     if (bmap_recptr->bmapctl_bufptr->dn_agstart != agstart) {
345         /* bad start idx at level holding an AG */
346         bmap_recptr->ctl_other_error = -1;
347         fsck_send_msg(fsck_BMAPCDMCSTI, 0);
348     }
349     if (bmap_recptr->bmapctl_bufptr->dn_agl2size != agg_recptr->log2_blkspersag) {
350         /* bad log2(fsblks per AG) */
351         bmap_recptr->ctl_other_error = -1;
352         fsck_send_msg(fsck_BMAPCL2BPAG, 0);
353     }
354     if (bmap_recptr->bmapctl_bufptr->dn_agsize != sb_ptr->s_agsize) {
355         /* bad fsblks per AG */
356         bmap_recptr->ctl_other_error = -1;
357         fsck_send_msg(fsck_BMAPCBPAG, 0);
358     }
359     if (bmap_recptr->bmapctl_bufptr->dn_maxfreebud != max_buddy) {

```

```

360         /* bad max free buddy system */
361         bmap_recptr->ctl_other_error = -1;
362         fsck_send_msg(fsck_BMAPCBMXB, 0);
363     }
364     for (agidx = 0; (agidx < MAXAG); agidx++) {
365         /* verify the free list */
366         if (bmap_recptr->bmpctl_bufptr->dn_agfree[agidx] != bmap_recptr->AGFree_tbl[agidx]) {
367             /* bad free blocks in the AG */
368             bmap_recptr->ctl_fctl_error = -1;
369             sprintf(message_parm_0, "%d", agidx);
370             msgprms[0] = message_parm_0;
371             msgprmidx[0] = 0;
372             fsck_send_msg(fsck_BMAPCAGNF, 1);
373         }
374     }
375 }
376 return (bcv_rc);
377 }
378
379 /*****
380  * NAME: dmap_pwmap_rebuild
381  *
382  * FUNCTION: Rebuild the pmap in the current block map dmap page.
383  *
384  * PARAMETERS:
385  *     pmap_freeblks - input - pointer to a variable in which the number
386  *                     of free blocks described by the dmap page
387  *                     will be returned.
388  *
389  * RETURNS:
390  *     success: FSCK_OK
391  *     failure: something else
392  */
393 int dmap_pwmap_rebuild(uint32_t * pmap_freeblks)
394 {
395     int bdpr_rc = FSCK_OK;
396     uint32_t bitmask;
397     int64_t wsp_pagenum;
398     uint32_t wsp_byteoffset;
399     struct fsck_blk_map_page *wsp_page;
400     uint32_t *wsp_bits = NULL;
401     int32_t map_wordidx, word_bitidx;
402     int8_t max_buddy;
403
404     /*
405      * locate the section of the workspace bit map which corresponds
406      * to this dmap's pmap
407      */
408     bdpr_rc = blkmap_find_bit(bmap_recptr->dmap_lstblk, &wsp_pagenum,
409                             &wsp_byteoffset, &bitmask);
410     if (bdpr_rc == FSCK_OK) {
411         bdpr_rc = blkmap_get_page(wsp_pagenum, &wsp_page);
412         if (bdpr_rc == FSCK_OK) {
413             wsp_bits =
414                 (uint32_t *) ((char *) wsp_page + wsp_byteoffset);
415         }
416     }
417
418     *pmap_freeblks = 0;
419
420     if (bdpr_rc == FSCK_OK) {
421         for (map_wordidx = 0; (map_wordidx < LPERDMAP); map_wordidx++) {
422             max_buddy =
423                 ujfs_maxbuddy((char *) &(wsp_bits[map_wordidx]));
424             bmap_recptr->dmap_wsp_sleafs[map_wordidx] = max_buddy;
425
426             bmap_recptr->dmap_bufptr->wmap[map_wordidx] =
427                 wsp_bits[map_wordidx];
428             /*
429              * copy the word from workspace to buffer,
430              * into both the working map and the permanent map.
431              */
432             bmap_recptr->dmap_bufptr->pmap[map_wordidx] = wsp_bits[map_wordidx];
433             /*
434              * count the free blocks described by the word
435              */
436             bitmask = 0x80000000u;
437             for (word_bitidx = 0; (word_bitidx < DBWORD);
438                 word_bitidx++) {
439
440                 if (!(wsp_bits[map_wordidx] & bitmask)) {
441                     /* it's free */
442                     (*pmap_freeblks)++;
443                     agg_recptr->free_blocks_in_aggregate++;
444                 } else {
445                     /* it's allocated */
446                     agg_recptr->blocks_used_in_aggregate++;
447                     /* end else it's allocated */
448                 }
449                 bitmask = bitmask >> 1; /* advance to the next bit */

```

```

450     }
451     }
452 }
453     return (bdpr_rc);
454 }
455
456 /*****
457  * NAME: dmap_pmap_verify
458  *
459  * FUNCTION: Verify the pmap in the current block map dmap page.
460  *
461  * PARAMETERS:
462  *     pmap_freeblks - input - pointer to a variable in which the number
463  *                     of free blocks described by the dmap page
464  *                     will be returned.
465  *
466  * RETURNS:
467  *     success: FSCK_OK
468  *     failure: something else
469  */
470 int dmap_pmap_verify(uint32_t * pmap_freeblks)
471 {
472     int bdpv_rc = FSCK_OK;
473     uint32_t bitmask;
474     int64_t wsp_pagenum;
475     uint32_t wsp_byteoffset;
476     struct fsck_blk_map_page *wsp_page;
477     uint32_t *wsp_bits = NULL;
478     int32_t map_wordidx, word_bitidx;
479     int8_t max_buddy;
480     uint32_t unmarked_range_first_ordno = 0;
481     int32_t unmarked_range_wordidx = 0;
482     int32_t unmarked_range_bitidx = 0;
483     int64_t size_of_unmarked_range = 0;
484     uint32_t marked_range_first_ordno = 0;
485     int32_t marked_range_wordidx = 0;
486     int32_t marked_range_bitidx = 0;
487     int64_t size_of_marked_range = 0;
488
489     /*
490      * locate the section of the workspace bit map which corresponds
491      * to this dmap's pmap
492      */
493     bdpv_rc = blkmap_find_bit(bmap_recptr->dmap_1stblk, &wsp_pagenum,
494                             &wsp_byteoffset, &bitmask);
495     if (bdpv_rc == FSCK_OK) {
496         bdpv_rc = blkmap_get_page(wsp_pagenum, &wsp_page);
497         if (bdpv_rc == FSCK_OK) { /* got the page */
498             wsp_bits =
499                 (uint32_t *) ((char *) wsp_page + wsp_byteoffset);
500         }
501     }
502
503     *pmap_freeblks = 0;
504
505     if (bdpv_rc == FSCK_OK) {
506         for (map_wordidx = 0; (map_wordidx < LPERDMAP); map_wordidx++) {
507             max_buddy =
508                 ujfs_maxbuddy((char *) &(wsp_bits[map_wordidx]));
509             bmap_recptr->dmap_wsp_sleafs[map_wordidx] = max_buddy;
510
511             bitmask = 0x80000000u;
512             for (word_bitidx = 0; (word_bitidx < DBWORD);
513                 word_bitidx++) {
514
515                 if (wsp_bits[map_wordidx] & bitmask) {
516                     agg_recptr->blocks_used_in_aggregate++;
517                     if (!(bmap_recptr->dmap_bufptr->pmap[map_wordidx] & bitmask)) {
518                         /* pmap says not */
519                         bmap_recptr->dmap_pmap_error =
520                             -1;
521
522                         if (size_of_unmarked_range == 0) {
523
524                             if (size_of_marked_range != 0) {
525                                 sprintf
526                                     (message_parm_0,
527                                      "%lld",
528                                       (long long)
529                                        size_of_marked_range);
530                                 msgprms[0] =
531                                     message_parm_0;
532                                 msgprmidx[0] =
533                                     0;
534                                 sprintf
535                                     (message_parm_1,
536                                      "%d",
537                                       marked_range_first_ordno);
538                                 msgprms[1] =
539                                     message_parm_1;

```

```

540         msgprmidx[1] =
541             0;
542         sprintf
543             (message_parm_2,
544              "%d",
545              marked_range_wordidx);
546         msgprms[2] =
547             message_parm_2;
548         msgprmidx[2] =
549             0;
550         sprintf
551             (message_parm_3,
552              "%d",
553              marked_range_bitidx);
554         msgprms[3] =
555             message_parm_3;
556         msgprmidx[3] =
557             0;
558         fsck_send_msg
559             (fsck_PMAPSBOFF,
560              4);
561         size_of_marked_range
562             = 0;
563     }
564     unmarked_range_first_ordno
565         =
566         bmap_recptr->
567         dmappg_ordno;
568     unmarked_range_wordidx =
569         map_wordidx;
570     unmarked_range_bitidx =
571         word_bitidx;
572     size_of_unmarked_range =
573         1;
574 } else {
575     /* not the first in the range */
576     size_of_unmarked_range++;
577 }
578 } else {
579     /* pmap agrees */
580     if (size_of_marked_range != 0) {
581         /* marked range ended */
582         sprintf(message_parm_0,
583                 "%lld",
584                 (long long)
585                 size_of_marked_range);
586         msgprms[0] =
587             message_parm_0;
588         msgprmidx[0] = 0;
589         sprintf(message_parm_1,
590                 "%d",
591                 marked_range_first_ordno);
592         msgprms[1] =
593             message_parm_1;
594         msgprmidx[1] = 0;
595         sprintf(message_parm_2,
596                 "%d",
597                 marked_range_wordidx);
598         msgprms[2] =
599             message_parm_2;
600         msgprmidx[2] = 0;
601         sprintf(message_parm_3,
602                 "%d",
603                 marked_range_bitidx);
604         msgprms[3] =
605             message_parm_3;
606         msgprmidx[3] = 0;
607         fsck_send_msg
608             (fsck_PMAPSBOFF, 4);
609         size_of_marked_range =
610             0;
611     }
612     if (size_of_unmarked_range != 0) {
613         /* unmarked range ended */
614         sprintf(message_parm_0,
615                 "%lld",
616                 (long long)
617                 size_of_unmarked_range);
618         msgprms[0] =
619             message_parm_0;
620         msgprmidx[0] = 0;
621         sprintf(message_parm_1,
622                 "%d",
623                 unmarked_range_first_ordno);
624         msgprmidx[1] = 0;
625         msgprms[1] =
626             message_parm_1;
627         sprintf(message_parm_2,
628                 "%d",
629                 unmarked_range_wordidx);

```



```

630         msgprmidx[2] = 0;
631         msgprms[2] =
632             message_parm_2;
633         sprintf(message_parm_3,
634             "%d",
635             unmarked_range_bitidx);
636         msgprmidx[3] = 0;
637         msgprms[3] =
638             message_parm_3;
639         fsck_send_msg
640             (fsck_PMAPSBON, 4);
641         size_of_unmarked_range =
642             0;
643     }
644 }
645 } else {
646     /* fsck says not allocated */
647     (*pmap_freeblks)++;
648     agg_recptr->free_blocks_in_aggregate++;
649
650     if ((bmap_recptr->dmap_bufptr->pmap[map_wordidx] & bitmask)) {
651         /* pmap says yes */
652         bmap_recptr->dmap_pmap_error =
653             -1;
654
655         if (size_of_marked_range == 0) {
656             if (size_of_unmarked_range != 0) {
657                 sprintf
658                     (message_parm_0,
659                     "%lld",
660                     (long long)
661                     size_of_unmarked_range);
662                 msgprms[0] =
663                     message_parm_0;
664                 msgprmidx[0] =
665                     0;
666                 sprintf
667                     (message_parm_1,
668                     "%d",
669                     unmarked_range_first_ordno);
670                 msgprms[1] =
671                     message_parm_1;
672                 msgprmidx[1] =
673                     0;
674                 sprintf
675                     (message_parm_2,
676                     "%d",
677                     unmarked_range_wordidx);
678                 msgprms[2] =
679                     message_parm_2;
680                 msgprmidx[2] =
681                     0;
682                 sprintf
683                     (message_parm_3,
684                     "%d",
685                     unmarked_range_bitidx);
686                 msgprms[3] =
687                     message_parm_3;
688                 msgprmidx[3] =
689                     0;
690                 fsck_send_msg
691                     (fsck_PMAPSBON,
692                     4);
693                 size_of_unmarked_range
694                     = 0;
695             }
696             marked_range_first_ordno
697                 =
698                 bmap_recptr->
699                 dmappg_ordno;
700             marked_range_wordidx =
701                 map_wordidx;
702             marked_range_bitidx =
703                 word_bitidx;
704             size_of_marked_range =
705                 1;
706         } else {
707             /* not the first in the range */
708             size_of_marked_range++;
709         }
710     } else {
711         /* pmap agrees */
712         if (size_of_marked_range != 0) {
713             sprintf(message_parm_0,
714                 "%lld",
715                 (long long)
716                 size_of_marked_range);
717             msgprms[0] =
718                 message_parm_0;
719             msgprmidx[0] = 0;

```

```

720         sprintf(message_parm_1,
721                 "%d",
722                 marked_range_first_ordno);
723     msgprms[1] =
724         message_parm_1;
725     msgprmidx[1] = 0;
726     sprintf(message_parm_2,
727             "%d",
728             marked_range_wordidx);
729     msgprms[2] =
730         message_parm_2;
731     msgprmidx[2] = 0;
732     sprintf(message_parm_3,
733             "%d",
734             marked_range_bitidx);
735     msgprms[3] =
736         message_parm_3;
737     msgprmidx[3] = 0;
738     fsck_send_msg
739         (fsck_PMAPSBOFF, 4);
740     size_of_marked_range =
741         0;
742     }
743     if (size_of_unmarked_range != 0) {
744         /* unmarked range ended */
745         sprintf(message_parm_0,
746                 "%lld",
747                 (long long)
748                 size_of_unmarked_range);
749         msgprms[0] =
750             message_parm_0;
751         msgprmidx[0] = 0;
752         sprintf(message_parm_1,
753                 "%d",
754                 unmarked_range_first_ordno);
755         msgprms[1] =
756             message_parm_1;
757         msgprmidx[1] = 0;
758         sprintf(message_parm_2,
759                 "%d",
760                 unmarked_range_wordidx);
761         msgprms[2] =
762             message_parm_2;
763         msgprmidx[2] = 0;
764         sprintf(message_parm_3,
765                 "%d",
766                 unmarked_range_bitidx);
767         msgprms[3] =
768             message_parm_3;
769         msgprmidx[3] = 0;
770         fsck_send_msg
771             (fsck_PMAPSBON, 4);
772         size_of_unmarked_range =
773             0;
774     }
775     }
776 }
777
778     /* advance to the next bit */
779     bitmask = bitmask >> 1;
780 }
781 }
782
783 if (size_of_marked_range != 0) {
784     /* marked range ended */
785     sprintf(message_parm_0, "%lld",
786             (long long) size_of_marked_range);
787     msgprms[0] = message_parm_0;
788     msgprmidx[0] = 0;
789     sprintf(message_parm_1, "%d", marked_range_first_ordno);
790     msgprms[1] = message_parm_1;
791     msgprmidx[1] = 0;
792     sprintf(message_parm_2, "%d", marked_range_wordidx);
793     msgprms[2] = message_parm_2;
794     msgprmidx[2] = 0;
795     sprintf(message_parm_3, "%d", marked_range_bitidx);
796     msgprms[3] = message_parm_3;
797     msgprmidx[3] = 0;
798     fsck_send_msg(fsck_PMAPSBOFF, 4);
799     size_of_marked_range = 0;
800 }
801 if (size_of_unmarked_range != 0) {
802     /* unmarked range ended */
803     sprintf(message_parm_0, "%lld",
804             (long long) size_of_unmarked_range);
805     msgprms[0] = message_parm_0;
806     msgprmidx[0] = 0;
807     sprintf(message_parm_1, "%d",
808             unmarked_range_first_ordno);
809     msgprms[1] = message_parm_1;

```

```

810         msgprmidx[1] = 0;
811         sprintf(message_parm_2, "%d", unmarked_range_wordidx);
812         msgprms[2] = message_parm_2;
813         msgprmidx[2] = 0;
814         sprintf(message_parm_3, "%d", unmarked_range_bitidx);
815         msgprms[3] = message_parm_3;
816         msgprmidx[3] = 0;
817         fsck_send_msg(fsck_PMAPSBON, 4);
818         size_of_unmarked_range = 0;
819     }
820 }
821     return (bdpv_rc);
822 }
823
824 /*****
825  * NAME: dmap_tree_rebuild
826  *
827  * FUNCTION: Rebuild the tree in the current block map dmap page.
828  *
829  * PARAMETERS:
830  *     root_data - input - pointer to a variable in which the data value
831  *                 stored in the root of the tree will be returned.
832  *
833  * RETURNS:
834  *     success: FSCK_OK
835  *     failure: something else
836  */
837 int dmap_tree_rebuild(int8_t * root_data)
838 {
839     int bdsr_rc = FSCK_OK;
840     struct fsck_stree_proc_parms stree_proc_parms;
841     struct fsck_stree_proc_parms *prms_ptr;
842
843     prms_ptr = &stree_proc_parms;
844
845 #define ppbt    prms_ptr->buf_tree
846
847     ppbt = (dmtree_t *) & (bmap_recptr->dmap_bufptr->tree);
848
849     ppbt->dmt_nleafs = prms_ptr->nleafs = LPERDMAP;
850     ppbt->dmt_l2nleafs = prms_ptr->l2nleafs = L2LPERDMAP;
851     ppbt->dmt_leafidx = prms_ptr->leafidx = LEAFIND;
852     ppbt->dmt_height = 4;
853     ppbt->dmt_budmin = prms_ptr->budmin = BUDMIN;
854
855     prms_ptr->buf_stree = &(ppbt->dmt_stree[0]);
856     prms_ptr->wsp_stree = bmap_recptr->dmap_wsp_stree;
857
858     bdsr_rc = stree_rebuild(prms_ptr, root_data);
859
860     return (bdsr_rc);
861 }
862
863 /*****
864  * NAME: dmap_tree_verify
865  *
866  * FUNCTION: Verify the tree in the current block map dmap page.
867  *
868  * PARAMETERS:
869  *     root_data - input - pointer to a variable in which the data value
870  *                 which should be stored in the root of the tree
871  *                 will be returned.
872  *
873  * RETURNS:
874  *     success: FSCK_OK
875  *     failure: something else
876  */
877 int dmap_tree_verify(int8_t * root_data)
878 {
879     int bdsv_rc = FSCK_OK;
880     int8_t tree_spec_errors = 0;
881     struct fsck_stree_proc_parms stree_proc_parms;
882     struct fsck_stree_proc_parms *prms_ptr;
883
884     prms_ptr = &stree_proc_parms;
885
886     prms_ptr->buf_tree = (dmtree_t *) & (bmap_recptr->dmap_bufptr->tree);
887
888     if (prms_ptr->buf_tree->dmt_nleafs != LPERDMAP) {
889         /* wrong number of leafs */
890         bmap_recptr->dmap_other_error = -1;
891         tree_spec_errors = -1;
892         msgprms[0] = message_parm_0;
893         msgprmidx[0] = fsck_dmap;
894         sprintf(message_parm_1, "%d", bmap_recptr->dmappg_ordno);
895         msgprms[1] = message_parm_1;
896         msgprmidx[1] = 0;
897         fsck_send_msg(fsck_BMAPBADNLF, 2);
898     }
899     if (prms_ptr->buf_tree->dmt_l2nleafs != L2LPERDMAP) {

```

```

900         /* wrong log2(nleafs) */
901         bmap_recptr->dmap_other_error = -1;
902         tree_spec_errors = -1;
903         msgprms[0] = message_parm_0;
904         msgprmidx[0] = fsck_dmap;
905         sprintf(message_parm_1, "%d", bmap_recptr->dmappg_ordno);
906         msgprms[1] = message_parm_1;
907         msgprmidx[1] = 0;
908         fsck_send_msg(fsck_BMAPBADL2NLF, 2);
909     }
910     if (prms_ptr->buf_tree->dm_t_leafidx != LEAFIND) {
911         /* wrong 1st leaf index */
912         bmap_recptr->dmap_other_error = -1;
913         tree_spec_errors = -1;
914         msgprms[0] = message_parm_0;
915         msgprmidx[0] = fsck_dmap;
916         sprintf(message_parm_1, "%d", bmap_recptr->dmappg_ordno);
917         msgprms[1] = message_parm_1;
918         msgprmidx[1] = 0;
919         fsck_send_msg(fsck_BMAPBADLFI, 2);
920     }
921     if (prms_ptr->buf_tree->dm_t_height != 4) {
922         /* wrong stree height */
923         bmap_recptr->dmap_other_error = -1;
924         tree_spec_errors = -1;
925         msgprms[0] = message_parm_0;
926         msgprmidx[0] = fsck_dmap;
927         sprintf(message_parm_1, "%d", bmap_recptr->dmappg_ordno);
928         msgprms[1] = message_parm_1;
929         msgprmidx[1] = 0;
930         fsck_send_msg(fsck_BMAPBADHT, 2);
931     }
932     if (prms_ptr->buf_tree->dm_t_budmin != BUDMIN) {
933         /* wrong min buddy value */
934         bmap_recptr->dmap_other_error = -1;
935         tree_spec_errors = -1;
936         msgprms[0] = message_parm_0;
937         msgprmidx[1] = fsck_dmap;
938         sprintf(message_parm_1, "%d", bmap_recptr->dmappg_ordno);
939         msgprms[1] = message_parm_1;
940         msgprmidx[2] = 0;
941         fsck_send_msg(fsck_BMAPBADBMN, 2);
942     }
943
944     /*
945     * if we found errors in the fields which specify the summary
946     * tree then we won't take the time to verify the tree itself.
947     *
948     * (The errors already detected would corrupt the summary tree,
949     *  so info about the bad tree would only be noise at this point.)
950     */
951     if (!tree_spec_errors) {
952         /* tree specification fields are ok */
953         prms_ptr->buf_stree = &(prms_ptr->buf_tree->dm_t_stree[0]);
954         prms_ptr->wsp_stree = bmap_recptr->dmappg_wsp_stree;
955         prms_ptr->nleafs = LPERDMAP;
956         prms_ptr->l2nleafs = L2LPERDMAP;
957         prms_ptr->leafidx = LEAFIND;
958         prms_ptr->budmin = BUDMIN;
959         prms_ptr->page_level = fsck_dmap;
960         prms_ptr->page_ordno = bmap_recptr->dmappg_ordno;
961         prms_ptr->lfval_error = &(bmap_recptr->dmappg_slfv_error);
962         prms_ptr->intval_error = &(bmap_recptr->dmappg_slv_error);
963
964         bdsrv_rc = stree_verify(prms_ptr, root_data);
965     }
966     return (bdsrv_rc);
967 }
968
969 /*****
970  * NAME: dmappg_rebuild
971  *
972  * FUNCTION: Rebuild the current dmap page.
973  *
974  * PARAMETERS:
975  *     stree_root_data - input - pointer to a variable in which to return
976  *                          the the data value in the root of the
977  *                          tree of the rebuilt page.
978  *
979  * RETURNS:
980  *     success: FSCK_OK
981  *     failure: something else
982  */
983 int dmappg_rebuild(int8_t * stree_root_data)
984 {
985     int bdsrv_rc = FSCK_OK;
986     uint32_t nblocks;
987     uint32_t dmap_freeblks;
988     uint32_t ag_num;
989

```

```

990     /*
991     * get the page to verify into the I/O buffer
992     */
993     bdmp_rc =
994         blktbl_dmap_get(bmap_recptr->dmap_lstblk,
995                        &(bmap_recptr->dmap_bufptr));
996
997     if (bdmp_rc == FSCK_OK) {
998         /* the page is in the buffer */
999         bmap_recptr->dmap_bufptr->start = bmap_recptr->dmap_lstblk;
1000         nblocks =
1001             MIN(BPERDMAP,
1002                (bmap_recptr->total_blocks - bmap_recptr->dmap_lstblk));
1003         bmap_recptr->dmap_bufptr->nblocks = nblocks;
1004         bdmp_rc = dmap_pwmap_rebuild(&dmap_freeblks);
1005         /*
1006          * subtract out any blocks which don't really exist but are
1007          * described by a dmap (phantom blocks always appear to be in use)
1008          */
1009         agg_recptr->blocks_used_in_aggregate = agg_recptr->blocks_used_in_aggregate - (BPERDMAP - nblocks)
1010     };
1011     if (bdmp_rc == FSCK_OK) {
1012         /* nothing strange during pmap reblid */
1013         bmap_recptr->free_blocks += dmap_freeblks;
1014         ag_num =
1015             bmap_recptr->dmap_lstblk >> agg_recptr->
1016             log2_blkspersag;
1017         bmap_recptr->AGFree_tbl[ag_num] += dmap_freeblks;
1018
1019         if (dmap_freeblks != nblocks) {
1020             /* not all of them are free */
1021             bmap_recptr->AGActive[ag_num] = -1;
1022         }
1023         bmap_recptr->dmap_bufptr->nfree = dmap_freeblks;
1024         bdmp_rc = dmap_tree_rebuild(stree_root_data);
1025         /*
1026          * write the page back to the device
1027          */
1028         if (bdmp_rc == FSCK_OK) {
1029             bdmp_rc =
1030                 blktbl_dmap_put(bmap_recptr->dmap_bufptr);
1031         }
1032     }
1033     return (bdmp_rc);
1034 }
1035
1036 /*****
1037  * NAME: dmappg_verify
1038  *
1039  * FUNCTION: Validate the current dmap page.
1040  *
1041  * PARAMETERS:
1042  *     stree_root_data - input - pointer to a variable in which to return
1043  *                          the the data value which should be stored
1044  *                          in the root of the tree of the page being
1045  *                          validated (and may also be the one stored
1046  *                          in the root of that tree)
1047  *
1048  * RETURNS:
1049  *     success: FSCK_OK
1050  *     failure: something else
1051  */
1052 int dmappg_verify(int8_t * stree_root_data)
1053 {
1054     int bdmpv_rc = FSCK_OK;
1055     uint32_t nblocks;
1056     uint32_t dmap_freeblks;
1057     uint32_t ag_num;
1058
1059     /*
1060     * get the page to verify into the I/O buffer
1061     */
1062     bdmpv_rc =
1063         blktbl_dmap_get(bmap_recptr->dmap_lstblk,
1064                        &(bmap_recptr->dmap_bufptr));
1065
1066     if (bdmpv_rc == FSCK_OK) {
1067         /* the page is in the buffer */
1068         if (bmap_recptr->dmap_bufptr->start != bmap_recptr->dmap_lstblk) {
1069             /* bad starting block */
1070             bmap_recptr->dmap_other_error = -1;
1071             sprintf(message_parm_0, "%d",
1072                    bmap_recptr->dmappg_ordno);
1073             msgprms[0] = message_parm_0;
1074             msgprmidx[0] = 0;
1075             fsck_send_msg(fsck_DMAPBADSTRT, 1);
1076         }
1077         nblocks =
1078             MIN(BPERDMAP,

```

```

1079         (bmap_recptr->total_blocks - bmap_recptr->dmap_1stblk));
1080     if (bmap_recptr->dmap_bufptr->nblocks != nblocks) {
1081         /* bad number of blocks */
1082         bmap_recptr->dmap_other_error = -1;
1083         sprintf(message_parm_0, "%d",
1084             bmap_recptr->dmappg_ordno);
1085         msgprms[0] = message_parm_0;
1086         msgprmidx[0] = 0;
1087         fsck_send_msg(fsck_DMAPBADNBLK, 1);
1088     }
1089     bdmpv_rc = dmap_pmap_verify(&dmap_freeblks);
1090     /*
1091      * subtract out any blocks which don't really exist but are described by a dmap
1092      * (phantom blocks always appear to be in use)
1093      */
1094     agg_recptr->blocks_used_in_aggregate = agg_recptr->blocks_used_in_aggregate - (BPERDMAP - nblocks
1095 );
1096     if (bdmpv_rc == FSCK_OK) {
1097         /* nothing strange during pmap ver */
1098         bmap_recptr->free_blocks += dmap_freeblks;
1099         ag_num =
1100             bmap_recptr->dmap_1stblk >> agg_recptr->
1101             log2_blkspersag;
1102         bmap_recptr->AGFree_tbl[ag_num] += dmap_freeblks;
1103         if (dmap_freeblks != nblocks) {
1104             /* not all of them are free */
1105             bmap_recptr->AGActive[ag_num] = -1;
1106         }
1107         if (bmap_recptr->dmap_bufptr->nfree != dmap_freeblks) {
1108             /* bad number of free blocks */
1109             bmap_recptr->dmap_other_error = -1;
1110             sprintf(message_parm_0, "%d",
1111                 bmap_recptr->dmappg_ordno);
1112             msgprms[0] = message_parm_0;
1113             msgprmidx[0] = 0;
1114             fsck_send_msg(fsck_DMAPBADNFREE, 1);
1115         }
1116         bdmpv_rc = dmap_tree_verify(stree_root_data);
1117     }
1118     return (bdmpv_rc);
1119 }
1120
1121 /*****
1122  * NAME:  init_bmap_info
1123  *
1124  * FUNCTION:  Initialize the bmap fsck global data area, pointed to by
1125  *            bmap_recptr, used to control JFS bmap processing.
1126  *
1127  * PARAMETERS:  none
1128  *
1129  * RETURNS:
1130  *     success:  FSCK_OK
1131  *     failure:  something else
1132  */
1133 int init_bmap_info()
1134 {
1135     int bii_rc = FSCK_OK;
1136     unsigned agidx;
1137
1138     memset(bmap_recptr, 0, sizeof (struct fsck_bmap_record));
1139
1140     memcpy((void *) &(bmap_recptr->eyecatcher), (void *) "bmaprecd", 8);
1141     memcpy((void *) &(bmap_recptr->bmpctlinf_eyecatcher),
1142         (void *) "bmapctl", 8);
1143     memcpy((void *) &(bmap_recptr->AGinf_eyecatcher), (void *) "AG info ",
1144         8);
1145     memcpy((void *) &(bmap_recptr->dmapinf_eyecatcher), (void *) "dmapinfo",
1146         8);
1147     memcpy((void *) &(bmap_recptr->L0inf_eyecatcher), (void *) "L0 info ",
1148         8);
1149     memcpy((void *) &(bmap_recptr->L1inf_eyecatcher), (void *) "L1 info ",
1150         8);
1151     memcpy((void *) &(bmap_recptr->L2inf_eyecatcher), (void *) "L2 info ",
1152         8);
1153
1154     bmap_recptr->bmpctl_bufptr =
1155         (struct dbmap *) agg_recptr->mapctl_buf_ptr;
1156     bmap_recptr->bmpctl_agg_fsblk_offset = BMAP_OFF / sb_ptr->s_bsize;
1157
1158     bmap_recptr->AGFree_tbl = &(agg_recptr->blkmap_wsp.AG_free[0]);
1159     bmap_recptr->dmap_wsp_stree =
1160         &(agg_recptr->blkmap_wsp.dmap_wsp_tree[0]);
1161     bmap_recptr->dmap_wsp_sleafs =
1162         &(agg_recptr->blkmap_wsp.dmap_wsp_leafs[0]);
1163     bmap_recptr->L0_wsp_stree = &(agg_recptr->blkmap_wsp.L0_wsp_tree[0]);
1164     bmap_recptr->L0_wsp_sleafs = &(agg_recptr->blkmap_wsp.L0_wsp_leafs[0]);
1165     bmap_recptr->L0_bufptr =
1166         (struct dmapctl *) &(agg_recptr->bmaplv_buf_ptr);
1167     bmap_recptr->L1_wsp_stree = &(agg_recptr->blkmap_wsp.L1_wsp_tree[0]);

```

```

1168     bmap_recptr->L1_wsp_sleafs = &(agg_recptr->blkmap_wsp.L1_wsp_sleafs[0]);
1169     bmap_recptr->L1_bufptr =
1170         (struct dmapctl *) &(agg_recptr->bmaplv_buf_ptr);
1171     bmap_recptr->L2_wsp_stree = &(agg_recptr->blkmap_wsp.L2_wsp_tree[0]);
1172     bmap_recptr->L2_wsp_sleafs = &(agg_recptr->blkmap_wsp.L2_wsp_sleafs[0]);
1173     bmap_recptr->L2_bufptr =
1174         (struct dmapctl *) &(agg_recptr->bmaplv_buf_ptr);
1175
1176     bmap_recptr->total_blocks =
1177         sb_ptr->s_size * sb_ptr->s_pbsize / sb_ptr->s_bsize;
1178     bmap_recptr->dmappg_count =
1179         (bmap_recptr->total_blocks + BPERDMAP - 1) / BPERDMAP;
1180     bmap_recptr->L0pg_count =
1181         (bmap_recptr->dmappg_count + LPERCTL - 1) / LPERCTL;
1182     if (bmap_recptr->L0pg_count > 1) {
1183         bmap_recptr->L1pg_count =
1184             (bmap_recptr->L0pg_count + LPERCTL - 1) / LPERCTL;
1185         if (bmap_recptr->L1pg_count > 1) {
1186             bmap_recptr->L2pg_count =
1187                 (bmap_recptr->L1pg_count + LPERCTL - 1) / LPERCTL;
1188         } else {
1189             bmap_recptr->L2pg_count = 0;
1190         }
1191     } else {
1192         bmap_recptr->L1pg_count = 0;
1193     }
1194
1195     bmap_recptr->free_blocks = 0;
1196     bmap_recptr->ctl_fctl_error = 0;
1197     bmap_recptr->ctl_other_error = 0;
1198
1199     for (agidx = 0; (agidx < MAXAG); agidx++) {
1200         bmap_recptr->AGActive[agidx] = 0;
1201     }
1202
1203     bmap_recptr->dmappg_ordno = bmap_recptr->L0pg_ordno = 0;
1204     bmap_recptr->L1pg_ordno = bmap_recptr->L2pg_1stblk = 0;
1205
1206     bmap_recptr->dmappg_idx = bmap_recptr->L0pg_idx = 0;
1207     bmap_recptr->L1pg_idx = 0;
1208
1209     bmap_recptr->dmap_1stblk = bmap_recptr->L0pg_1stblk = 0;
1210     bmap_recptr->L1pg_1stblk = bmap_recptr->L2pg_1stblk = 0;
1211
1212     bmap_recptr->dmap_pmap_error = 0;
1213
1214     bmap_recptr->dmap_slfv_error = bmap_recptr->L0pg_slfv_error = 0;
1215     bmap_recptr->L1pg_slfv_error = bmap_recptr->L2pg_slfv_error = 0;
1216
1217     bmap_recptr->dmap_slnv_error = bmap_recptr->L0pg_slnv_error = 0;
1218     bmap_recptr->L1pg_slnv_error = bmap_recptr->L2pg_slnv_error = 0;
1219
1220     bmap_recptr->dmap_other_error = bmap_recptr->L0pg_other_error = 0;
1221     bmap_recptr->L1pg_other_error = bmap_recptr->L2pg_other_error = 0;
1222
1223     return (bii_rc);
1224 }
1225
1226 /*****
1227  * NAME: Ln_tree_rebuild
1228  *
1229  * FUNCTION: Rebuild the tree in the current summary page.
1230  *
1231  * PARAMETERS:
1232  *     level           - input - Summary level of the bmap summary page
1233  *                     containing the tree to rebuild
1234  *     first_blk       - input - First aggregate block described by the bmap
1235  *                     summary page containing the tree to rebuild
1236  *     addr_buf_ptr    - input - pointer to a variable in which to return the
1237  *                     address of the buffer in which the page has
1238  *                     been rebuilt (and then written to the
1239  *                     aggregate)
1240  *     root_data       - input - pointer to a variable in which to return the
1241  *                     data value in the root of the rebuilt tree.
1242  *
1243  * RETURNS:
1244  *     success: FSCK_OK
1245  *     failure: something else
1246  */
1247 int Ln_tree_rebuild(int level, int64_t first_blk, struct dmapctl **addr_buf_ptr,
1248                    int8_t * root_data)
1249 {
1250     int blsr_rc = FSCK_OK;
1251     struct fsck_stree_proc_parms stree_proc_parms;
1252     struct fsck_stree_proc_parms *prms_ptr;
1253
1254     /*
1255     * get the page to verify into the I/O buffer
1256     */
1257     blsr_rc = blktbl_Ln_page_get(level, first_blk, addr_buf_ptr);

```

```

1258     if (blsr_rc == FSCK_OK) {
1259         /* the page is in the buffer */
1260         prms_ptr = &stree_proc_parms;
1261         /* these are just big trees */
1262         prms_ptr->buf_tree = (dmtree_t *) * addr_buf_ptr;
1263
1264         switch (level) {
1265             case 0:{
1266                 prms_ptr->wsp_stree = bmap_recptr->L0_wsp_stree;
1267                 break;
1268             }
1269             case 1:{
1270                 prms_ptr->wsp_stree = bmap_recptr->L1_wsp_stree;
1271                 break;
1272             }
1273             default:{
1274                 prms_ptr->wsp_stree = bmap_recptr->L2_wsp_stree;
1275                 break;
1276             }
1277         }
1278
1279         prms_ptr->buf_tree->dm_t_nleafs = prms_ptr->nleafs = LPERCTL;
1280         prms_ptr->buf_tree->dm_t_l2nleafs = prms_ptr->l2nleafs =
1281             L2LPERCTL;
1282         prms_ptr->buf_tree->dm_t_leafidx = prms_ptr->leafidx =
1283             CTLLLEAFIND;
1284         prms_ptr->buf_tree->dm_t_height = 5;
1285         prms_ptr->buf_tree->dm_t_budmin = L2BPERDMAP + level * L2LPERCTL;
1286         prms_ptr->budmin = prms_ptr->buf_tree->dm_t_budmin;
1287
1288         prms_ptr->buf_stree = &(prms_ptr->buf_tree->dm_t_stree[0]);
1289
1290         blsr_rc = stree_rebuild(prms_ptr, root_data);
1291
1292         /*
1293          * put the page back into the file
1294          */
1295         if (blsr_rc == FSCK_OK) {
1296             blsr_rc = blk_tbl_Ln_page_put(*addr_buf_ptr);
1297         }
1298     }
1299     return (blsr_rc);
1300 }
1301
1302
1303 /*****
1304  * NAME: Ln_tree_verify
1305  *
1306  * FUNCTION: Verify the tree in the current summary page.
1307  *
1308  * PARAMETERS:
1309  *     level          - input - Summary level of the bmap summary page
1310  *                    containing the tree to verify
1311  *     first_blk      - input - First aggregate block described by the bmap
1312  *                    summary page containing the tree to verify
1313  *     addr_buf_ptr   - input - pointer to a variable in which to return the
1314  *                    address of the buffer into which the page has
1315  *                    been read so that its contents can be verified
1316  *     root_data      - input - pointer to a variable in which to return the
1317  *                    data value which should be in the root of the
1318  *                    tree....and may actually be stored there.
1319  *
1320  * RETURNS:
1321  *     success: FSCK_OK
1322  *     failure: something else
1323  */
1324 int Ln_tree_verify(int level, int64_t first_blk, struct dmapctl **addr_buf_ptr,
1325                  int8_t * root_data)
1326 {
1327     int blsv_rc = FSCK_OK;
1328     int8_t *other_errors;
1329     int8_t tree_spec_errors = 0;
1330     struct fsck_stree_proc_parms stree_proc_parms;
1331     struct fsck_stree_proc_parms *prms_ptr;
1332
1333     /*
1334      * get the page to verify into the I/O buffer
1335      */
1336     blsv_rc = blk_tbl_Ln_page_get(level, first_blk, addr_buf_ptr);
1337
1338     if (blsv_rc == FSCK_OK) {
1339         /* the page is in the buffer */
1340         prms_ptr = &stree_proc_parms;
1341         /* these are just big trees */
1342         prms_ptr->buf_tree = (dmtree_t *) * addr_buf_ptr;
1343
1344         switch (level) {
1345             case 0:{
1346                 prms_ptr->wsp_stree = bmap_recptr->L0_wsp_stree;
1347                 prms_ptr->page_ordno = bmap_recptr->L0pg_ordno;

```



```

1348         prms_ptr->lfval_error =
1349             &(bmap_recptr->L0pg_slfv_error);
1350         prms_ptr->intval_error =
1351             &(bmap_recptr->L0pg_slnv_error);
1352         prms_ptr->page_level = fsck_L0;
1353         other_errors = &(bmap_recptr->L0pg_other_error);
1354         break;
1355     }
1356     case 1: {
1357         prms_ptr->wsp_stree = bmap_recptr->L1_wsp_stree;
1358         prms_ptr->page_ordno = bmap_recptr->L1pg_ordno;
1359         prms_ptr->lfval_error =
1360             &(bmap_recptr->L1pg_slfv_error);
1361         prms_ptr->intval_error =
1362             &(bmap_recptr->L1pg_slnv_error);
1363         prms_ptr->page_level = fsck_L1;
1364         other_errors = &(bmap_recptr->L1pg_other_error);
1365         break;
1366     }
1367     default: {
1368         prms_ptr->wsp_stree = bmap_recptr->L2_wsp_stree;
1369         prms_ptr->page_ordno = 0;
1370         prms_ptr->lfval_error =
1371             &(bmap_recptr->L2pg_slfv_error);
1372         prms_ptr->intval_error =
1373             &(bmap_recptr->L2pg_slnv_error);
1374         prms_ptr->page_level = fsck_L2;
1375         other_errors = &(bmap_recptr->L2pg_other_error);
1376         break;
1377     }
1378 }
1379
1380 if (prms_ptr->buf_tree->dmtd_nleafs != LPERCTL) {
1381     /* wrong number of leafs */
1382     *other_errors = -1;
1383     tree_spec_errors = -1;
1384     msgprms[0] = message_parm_0;
1385     msgprmidx[0] = prms_ptr->page_level;
1386     sprintf(message_parm_1, "%d", prms_ptr->page_ordno);
1387     msgprms[1] = message_parm_1;
1388     msgprmidx[1] = 0;
1389     fsck_send_msg(fsck_BMAPBADNLF, 2);
1390 }
1391 if (prms_ptr->buf_tree->dmtd_l2nleafs != L2LPERCTL) {
1392     /* wrong log2(nleafs) */
1393     *other_errors = -1;
1394     tree_spec_errors = -1;
1395     msgprms[0] = message_parm_0;
1396     msgprmidx[0] = prms_ptr->page_level;
1397     sprintf(message_parm_1, "%d", prms_ptr->page_ordno);
1398     msgprms[1] = message_parm_1;
1399     msgprmidx[1] = 0;
1400     fsck_send_msg(fsck_BMAPBADL2NLF, 2);
1401 }
1402 if (prms_ptr->buf_tree->dmtd_leafidx != CTLLEAFIND) {
1403     /* wrong 1st leaf index */
1404     *other_errors = -1;
1405     tree_spec_errors = -1;
1406     msgprms[0] = message_parm_0;
1407     msgprmidx[0] = prms_ptr->page_level;
1408     sprintf(message_parm_1, "%d", prms_ptr->page_ordno);
1409     msgprms[1] = message_parm_1;
1410     msgprmidx[1] = 0;
1411     fsck_send_msg(fsck_BMAPBADLFI, 2);
1412 }
1413 if (prms_ptr->buf_tree->dmtd_height != 5) {
1414     /* wrong stree height */
1415     *other_errors = -1;
1416     tree_spec_errors = -1;
1417     msgprms[0] = message_parm_0;
1418     msgprmidx[0] = prms_ptr->page_level;
1419     sprintf(message_parm_1, "%d", prms_ptr->page_ordno);
1420     msgprms[1] = message_parm_1;
1421     msgprmidx[1] = 0;
1422     fsck_send_msg(fsck_BMAPBADHT, 2);
1423 }
1424 if (prms_ptr->buf_tree->dmtd_budmin != (L2BPERDMAP + level * L2LPERCTL)) {
1425     /* wrong min buddy value */
1426     *other_errors = -1;
1427     tree_spec_errors = -1;
1428     msgprms[0] = message_parm_0;
1429     msgprmidx[0] = prms_ptr->page_level;
1430     sprintf(message_parm_1, "%d", prms_ptr->page_ordno);
1431     msgprms[1] = message_parm_1;
1432     msgprmidx[1] = 0;
1433     fsck_send_msg(fsck_BMAPBADBMN, 2);
1434 }
1435
1436 /*
1437  * if we found errors in the fields which specify the summary

```

```

1438         * tree then we won't take the time to verify the tree itself.
1439         *
1440         * (The errors already detected would corrupt the summary tree,
1441         *  so info about the bad tree would only be noise at this point.)
1442         */
1443         if (!tree_spec_errors) {
1444             /* tree specification fields are ok */
1445             prms_ptr->buf_stree =
1446                 &(prms_ptr->buf_tree->dm_t_stree[0]);
1447             prms_ptr->nleafs = LPERCTL;
1448             prms_ptr->budmin = (L2BPERDMAP + level * L2LPERCTL);
1449             prms_ptr->l2nleafs = L2LPERCTL;
1450             prms_ptr->leafidx = CTLLLEAFIND;
1451
1452             blsv_rc = stree_verify(prms_ptr, root_data);
1453
1454         }
1455     }
1456     return (blsv_rc);
1457 }
1458
1459 /*****
1460  * NAME: rebuild_blkall_map
1461  *
1462  * FUNCTION: Rebuild the JFS Aggregate Block Map in the aggregate.
1463  *
1464  * PARAMETERS: none
1465  *
1466  * RETURNS:
1467  *     success: FSCK_OK
1468  *     failure: something else
1469  */
1470 int rebuild_blkall_map()
1471 {
1472     int rbam_rc = FSCK_OK;
1473     int8_t sumtree_root_data;
1474     uint32_t leafidx;
1475
1476     #define MAXIDX (LPERCTL - 1)
1477
1478     rbam_rc = init_bmap_info();
1479
1480     /*
1481      * since the dmap I/O buffer is really the same storage as the
1482      * IAG I/O buffer, flush out any pending writes that may remain
1483      * from IAG processing.
1484      */
1485     rbam_rc = iags_flush();
1486
1487     /*
1488      * rebuild the dmap pages. Rebuild each L0 and L1 page
1489      * if and when the information for is complete.
1490      */
1491     while ((rbam_rc == FSCK_OK)
1492            && (bmap_recptr->dmappg_ordno < bmap_recptr->dmappg_count)) {
1493
1494         rbam_rc = dmappg_rebuild(&sumtree_root_data);
1495         if (rbam_rc == FSCK_OK) {
1496             /*
1497              * the data in the dmap summary tree root goes into a leaf of
1498              * the current L0 page
1499              */
1500             bmap_recptr->L0_wsp_sleafs[bmap_recptr->dmappg_idx] =
1501                 sumtree_root_data;
1502
1503             /* move to next dmap page */
1504             bmap_recptr->dmappg_ordno++;
1505             bmap_recptr->dmap_lstblk += BPERDMAP;
1506
1507             if (bmap_recptr->dmappg_idx < MAXIDX) {
1508                 /* still gathering info about this L0 page */
1509                 bmap_recptr->dmappg_idx++;
1510             } else {
1511                 /* we have all the info needed for the current L0 page */
1512                 bmap_recptr->dmappg_idx = 0;
1513                 rbam_rc =
1514                     Ln_tree_rebuild(0, bmap_recptr->L0pg_lstblk,
1515                                     &(bmap_recptr->L0_bufptr),
1516                                     &sumtree_root_data);
1517
1518                 if (rbam_rc == FSCK_OK) {
1519                     /*
1520                      * the data in the L0 summary tree root goes into
1521                      * a leaf of the current L1 page
1522                      */
1523                     bmap_recptr->L1_wsp_sleafs[bmap_recptr->
1524                                             L0pg_idx] =
1525                         sumtree_root_data;
1526                     /* move to the next L0 page */
1527                     bmap_recptr->L0pg_ordno++;

```

```

1528         bmap_recptr->L0pg_lstblk =
1529             bmap_recptr->dmap_lstblk;
1530
1531         if (bmap_recptr->L0pg_idx < MAXIDX) {
1532             /* still gathering info about this L1 page */
1533             bmap_recptr->L0pg_idx++;
1534         } else {
1535             /* we have all the info needed for the current L1 page */
1536             bmap_recptr->L0pg_idx = 0;
1537             rbam_rc =
1538                 Ln_tree_rebuild(1,
1539                                bmap_recptr->
1540                                L1pg_lstblk,
1541                                &
1542                                (bmap_recptr->
1543                                L1_bufptr),
1544                                &sumtree_root_data);
1545
1546             if (rbam_rc == FSCK_OK) {
1547                 /*
1548                  * the data in the L1 summary tree root goes into
1549                  * a leaf of the current L2 page
1550                  */
1551                 bmap_recptr->
1552                     L2_wsp_sleafs
1553                     [bmap_recptr->
1554                     L1pg_idx] =
1555                     sumtree_root_data;
1556
1557                 /* move to the next L1 page */
1558                 bmap_recptr->L1pg_ordno++;
1559                 bmap_recptr->
1560                     L1pg_lstblk =
1561                     bmap_recptr->
1562                     dmap_lstblk;
1563
1564                 /*
1565                  * note that there is always AT MOST a single L2 page
1566                  */
1567                 bmap_recptr->L1pg_idx++;
1568             }
1569         }
1570     }
1571 }
1572 }
1573 }
1574 }
1575
1576 /*
1577  * finish up the partial pages
1578  */
1579
1580 if (rbam_rc == FSCK_OK) {
1581     if (bmap_recptr->dmappg_idx != 0) {
1582         /* there's a partial L0 page */
1583         for (leafidx = bmap_recptr->dmappg_idx;
1584              (leafidx <= MAXIDX); leafidx++) {
1585             bmap_recptr->L0_wsp_sleafs[leafidx] = -1;
1586         }
1587         rbam_rc = Ln_tree_rebuild(0, bmap_recptr->L0pg_lstblk,
1588                                   &(bmap_recptr->L0_bufptr),
1589                                   &sumtree_root_data);
1590         if (rbam_rc == FSCK_OK) {
1591             /*
1592              * the data in the L0 summary tree root goes into
1593              * a leaf of the current L1 page
1594              */
1595             bmap_recptr->L1_wsp_sleafs[bmap_recptr->
1596                                       L0pg_idx] =
1597                 sumtree_root_data;
1598             bmap_recptr->L0pg_idx++;
1599         }
1600     }
1601 }
1602 if (rbam_rc == FSCK_OK) {
1603     if ((bmap_recptr->L1pg_count > 0) && (bmap_recptr->L0pg_idx != 0)) {
1604         /* there's enough data for an L1 level, and there's a partial L1 page */
1605         for (leafidx = bmap_recptr->L0pg_idx;
1606              (leafidx <= MAXIDX); leafidx++) {
1607             bmap_recptr->L1_wsp_sleafs[leafidx] = -1;
1608         }
1609         rbam_rc = Ln_tree_rebuild(1, bmap_recptr->L1pg_lstblk,
1610                                   &(bmap_recptr->L1_bufptr),
1611                                   &sumtree_root_data);
1612         if (rbam_rc == FSCK_OK) {
1613             /*
1614              * the data in the L0 summary tree root goes into
1615              * a leaf of the current L1 page
1616              */
1617             bmap_recptr->L2_wsp_sleafs[bmap_recptr->

```

```

1618                                     L1pg_idx] =
1619                                     sumtree_root_data;
1620                                     bmap_recptr->L1pg_idx++;
1621                                     }
1622                                     }
1623                                     }
1624                                     if (rbam_rc == FSCK_OK) {
1625                                         if ((bmap_recptr->L2pg_count > 0) && (bmap_recptr->L1pg_idx != 0)) {
1626                                             /* there's enough data for an L2 level, and there's a partial L2 page */
1627                                             for (leafidx = bmap_recptr->L1pg_idx;
1628                                                  (leafidx <= MAXIDX); leafidx++) {
1629                                                 bmap_recptr->L2_wsp_sleafs[leafidx] = -1;
1630                                             }
1631                                             rbam_rc = Ln_tree_rebuild(2, bmap_recptr->L2pg_1stblk,
1632                                                                    &(bmap_recptr->L2_bufptr),
1633                                                                    &sumtree_root_data);
1634                                         }
1635                                     }
1636                                     /*
1637                                     * Now go verify the Block Allocation Map Control page
1638                                     */
1639                                     if (rbam_rc == FSCK_OK) {
1640                                         rbam_rc = ctlpage_rebuild(sumtree_root_data);
1641                                     }
1642                                     return (rbam_rc);
1643                                     }
1644                                     }
1645                                     /*****
1646                                     * NAME: stree_rebuild
1647                                     *
1648                                     * FUNCTION: Rebuild the specified summary tree.
1649                                     *
1650                                     * PARAMETERS:
1651                                     *     prms_ptr    - input - pointer to a data area describing the tree
1652                                     *                -          to rebuild and containing the dmap which the
1653                                     *                -          tree summarizes.
1654                                     *     root_data  - input - pointer to a variable in which to return the
1655                                     *                -          data value stored in the root node of the tree
1656                                     *
1657                                     * RETURNS:
1658                                     *     success: FSCK_OK
1659                                     *     failure: something else
1660                                     */
1661                                     int stree_rebuild(struct fsck_stree_proc_parms *prms_ptr, int8_t * root_data)
1662                                     {
1663                                         int bsr_rc = FSCK_OK;
1664                                         uint32_t node_idx, last_leaf_idx;
1665
1666                                         /*
1667                                         * copy the leaf data into the buffer
1668                                         */
1669                                         last_leaf_idx = prms_ptr->leafidx + prms_ptr->nleafs - 1;
1670                                         for (node_idx = prms_ptr->leafidx; (node_idx <= last_leaf_idx);
1671                                              node_idx++) {
1672                                             prms_ptr->buf_stree[node_idx] = prms_ptr->wsp_stree[node_idx];
1673                                         }
1674
1675                                         /*
1676                                         * build the summary tree from the "raw" leaf values
1677                                         */
1678                                         *root_data =
1679                                         ujfs_adjtree(prms_ptr->buf_stree, prms_ptr->l2nleafs,
1680                                                    prms_ptr->budmin);
1681
1682                                         return (bsr_rc);
1683                                     }
1684                                     /*****
1685                                     * NAME: stree_verify
1686                                     *
1687                                     * FUNCTION: Verify the specified summary tree.
1688                                     *
1689                                     * PARAMETERS:
1690                                     *     prms_ptr    - input - pointer to a data area describing the tree
1691                                     *                -          to verify and containing the dmap which the
1692                                     *                -          tree summarizes.
1693                                     *     root_data  - input - pointer to a variable in which to return the
1694                                     *                -          data value which should be in the root node
1695                                     *                -          of the tree (and may in fact be in the root
1696                                     *                -          node of the tree)
1697                                     *
1698                                     * RETURNS:
1699                                     *     success: FSCK_OK
1700                                     *     failure: something else
1701                                     */
1702                                     int stree_verify(struct fsck_stree_proc_parms *prms_ptr, int8_t * root_data)
1703                                     {
1704                                         int bsv_rc = FSCK_OK;
1705                                         uint32_t node_idx, last_leaf_idx;
1706                                     }
1707                                     }

```

```

1708     /*
1709     * build the summary tree from the "raw" leaf values
1710     */
1711     *root_data =
1712         ujfs_adjtree(prms_ptr->wsp_stree, prms_ptr->l2nleafs,
1713                     prms_ptr->budmin);
1714
1715     /*
1716     * Now see if the tree in the buffer matches the one we just
1717     * built in the workspace.
1718     *
1719     * We distinguish between incorrect internal nodes and incorrect
1720     * leaf nodes because they can be symptoms of different problems.
1721     */
1722     for (node_idx = 0; (node_idx < prms_ptr->leafidx); node_idx++) {
1723         if (prms_ptr->buf_stree[node_idx] != prms_ptr->wsp_stree[node_idx]) {
1724             /* they don't match! */
1725             *(prms_ptr->intval_error) = -1;
1726             sprintf(message_parm_0, "%d", node_idx);
1727             msgprms[0] = message_parm_0;
1728             msgprmidx[0] = 0;
1729             msgprms[1] = message_parm_1;
1730             msgprmidx[1] = prms_ptr->page_level;
1731             sprintf(message_parm_2, "%d", prms_ptr->page_ordno);
1732             msgprms[2] = message_parm_2;
1733             msgprmidx[2] = 0;
1734             fsck_send_msg(fsck_BMAPBADLNV, 3);
1735         }
1736     }
1737
1738     last_leaf_idx = prms_ptr->leafidx + prms_ptr->nleafs - 1;
1739     for (node_idx = prms_ptr->leafidx; (node_idx <= last_leaf_idx);
1740         node_idx++) {
1741         if (prms_ptr->buf_stree[node_idx] != prms_ptr->wsp_stree[node_idx]) {
1742             /* they don't match! */
1743             *(prms_ptr->lfval_error) = -1;
1744             sprintf(message_parm_0, "%d", node_idx);
1745             msgprms[0] = message_parm_0;
1746             msgprmidx[0] = 0;
1747             msgprms[1] = message_parm_1;
1748             msgprmidx[1] = prms_ptr->page_level;
1749             sprintf(message_parm_2, "%d", prms_ptr->page_ordno);
1750             msgprms[2] = message_parm_2;
1751             msgprmidx[2] = 0;
1752             fsck_send_msg(fsck_BMAPBADLFV, 3);
1753         }
1754     }
1755
1756     return (bsv_rc);
1757 }
1758
1759 /*****
1760 * NAME: verify_blkall_map
1761 *
1762 * FUNCTION: Validate the JFS Aggregate Block Map for the aggregate.
1763 *
1764 * PARAMETERS: none
1765 *
1766 * RETURNS:
1767 *     success: FSCK_OK
1768 *     failure: something else
1769 */
1770 int verify_blkall_map()
1771 {
1772     int vbam_rc = FSCK_OK;
1773     int8_t sumtree_root_data;
1774     uint32_t leafidx;
1775
1776     #define MAXIDX (LPERCTL - 1)
1777
1778     vbam_rc = init_bmap_info();
1779
1780     /*
1781     * since the dmap I/O buffer is really the same storage as the
1782     * IAG I/O buffer, flush out any pending writes that may remain
1783     * from IAG processing.
1784     */
1785     vbam_rc = iags_flush();
1786
1787     /*
1788     * Verify the dmap pages. Verify each L0 and L1 page
1789     * if and when the information for is complete.
1790     */
1791     while ((vbam_rc == FSCK_OK)
1792           && (bmap_recptr->dmappg_ordno < bmap_recptr->dmappg_count)) {
1793
1794         vbam_rc = dmappg_verify(&sumtree_root_data);
1795         if (vbam_rc == FSCK_OK) {
1796             /*
1797              * the data in the dmap summary tree root goes into a leaf of

```

```

1798         * the current L0 page
1799         */
1800         bmap_recptr->L0_wsp_sleafs[bmap_recptr->dmappg_idx] =
1801             sumtree_root_data;
1802         /* move to next dmap page */
1803         bmap_recptr->dmappg_ordno++;
1804         bmap_recptr->dmap_1stblk += BPERDMAP;
1805
1806         if (bmap_recptr->dmappg_idx < MAXIDX) {
1807             /* still gathering info about this L0 page */
1808             bmap_recptr->dmappg_idx++;
1809         } else {
1810             /* we have all the info needed for the current L0 page */
1811             bmap_recptr->dmappg_idx = 0;
1812             vbam_rc =
1813                 Ln_tree_verify(0, bmap_recptr->L0pg_1stblk,
1814                               &(bmap_recptr->L0_bufptr),
1815                               &sumtree_root_data);
1816
1817             if (vbam_rc == FSCK_OK) {
1818                 /*
1819                  * the data in the L0 summary tree root goes into
1820                  * a leaf of the current L1 page
1821                  */
1822                 bmap_recptr->L1_wsp_sleafs[bmap_recptr->
1823                                         L0pg_idx] =
1824                     sumtree_root_data;
1825                 /* move to the next L0 page */
1826                 bmap_recptr->L0pg_ordno++;
1827                 bmap_recptr->L0pg_1stblk =
1828                     bmap_recptr->dmap_1stblk;
1829                 if (bmap_recptr->L0pg_idx < MAXIDX) {
1830                     /* still gathering info about this L1 page */
1831                     bmap_recptr->L0pg_idx++;
1832                 } else {
1833                     /* we have all the info needed for the current L1 page */
1834                     bmap_recptr->L0pg_idx = 0;
1835                     vbam_rc =
1836                         Ln_tree_verify(1,
1837                                       bmap_recptr->
1838                                           L1pg_1stblk,
1839                                       &
1840                                           (bmap_recptr->
1841                                               L1_bufptr),
1842                                       &sumtree_root_data);
1843                     if (vbam_rc == FSCK_OK) {
1844                         /*
1845                          * the data in the L1 summary tree root goes into
1846                          * a leaf of the current L2 page
1847                          */
1848                         bmap_recptr->
1849                             L2_wsp_sleafs
1850                             [bmap_recptr->
1851                                 L1pg_idx] =
1852                             sumtree_root_data;
1853                         /* move to the next L1 page */
1854                         bmap_recptr->L1pg_ordno++;
1855                         bmap_recptr->
1856                             L1pg_1stblk =
1857                             bmap_recptr->
1858                                 dmap_1stblk;
1859                         /*
1860                          * note that there is always AT MOST a single L2 page
1861                          */
1862                         bmap_recptr->L1pg_idx++;
1863                     }
1864                 }
1865             }
1866         }
1867     }
1868 }
1869 /*
1870  * finish up the partial pages
1871  */
1872 if (vbam_rc == FSCK_OK) {
1873     if (bmap_recptr->dmappg_idx != 0) {
1874         /* there's a partial L0 page */
1875         for (leafidx = bmap_recptr->dmappg_idx;
1876             (leafidx <= MAXIDX); leafidx++) {
1877             bmap_recptr->L0_wsp_sleafs[leafidx] = -1;
1878         }
1879         vbam_rc = Ln_tree_verify(0, bmap_recptr->L0pg_1stblk,
1880                                 &(bmap_recptr->L0_bufptr),
1881                                 &sumtree_root_data);
1882         if (vbam_rc == FSCK_OK) {
1883             /*
1884              * the data in the L0 summary tree root goes into
1885              * a leaf of the current L1 page
1886              */
1887             bmap_recptr->L1_wsp_sleafs[bmap_recptr->

```

```

1888                                     L0pg_idx] =
1889                                     sumtree_root_data;
1890                                     bmap_recptr->L0pg_idx++;
1891                                 }
1892                             }
1893                         }
1894                     if (vbam_rc == FSCK_OK) {
1895                         if ((bmap_recptr->L1pg_count > 0) && (bmap_recptr->L0pg_idx != 0)) {
1896                             /* there's enough data for an L1 level, and there's a partial L1 page */
1897                             for (leafidx = bmap_recptr->L0pg_idx;
1898                                  (leafidx <= MAXIDX); leafidx++) {
1899                                 bmap_recptr->L1_wsp_sleafs[leafidx] = -1;
1900                             }
1901                             /* end for leafidx */
1902
1903                             vbam_rc = Ln_tree_verify(1, bmap_recptr->L1pg_1stblk,
1904                                                       &(bmap_recptr->L1_bufptr),
1905                                                       &sumtree_root_data);
1906
1907                             if (vbam_rc == FSCK_OK) {
1908                                 /*
1909                                  * the data in the L0 summary tree root goes into
1910                                  * a leaf of the current L1 page
1911                                  */
1912                                 bmap_recptr->L2_wsp_sleafs[bmap_recptr->
1913                                                         L1pg_idx] =
1914                                     sumtree_root_data;
1915                                 bmap_recptr->L1pg_idx++;
1916                             }
1917                         }
1918                     if (vbam_rc == FSCK_OK) {
1919                         if ((bmap_recptr->L2pg_count > 0) && (bmap_recptr->L1pg_idx != 0)) {
1920                             /* there's enough data for an L2 level, and there's a partial L2 page */
1921                             for (leafidx = bmap_recptr->L1pg_idx;
1922                                  (leafidx <= MAXIDX); leafidx++) {
1923                                 bmap_recptr->L2_wsp_sleafs[leafidx] = -1;
1924                             }
1925
1926                             vbam_rc = Ln_tree_verify(2, bmap_recptr->L2pg_1stblk,
1927                                                       &(bmap_recptr->L2_bufptr),
1928                                                       &sumtree_root_data);
1929                         }
1930                     }
1931                     /*
1932                      * Now go verify the Block Allocation Map Control page
1933                      */
1934                     if (vbam_rc == FSCK_OK) {
1935                         vbam_rc = ctlpage_verify(sumtree_root_data);
1936                     }
1937                     /*
1938                      * issue summary messages about the Block Allocation Map validation
1939                      */
1940                     if (vbam_rc == FSCK_OK) {
1941                         vbam_rc = verify_blkall_summary_msgs();
1942                     }
1943                     return (vbam_rc);
1944                 }
1945
1946 /*****
1947  * NAME: verify_blkall_summary_msgs
1948  *
1949  * FUNCTION: Issue summary messages with the results of JFS Aggregate Block
1950  *           Map validation.
1951  *
1952  * PARAMETERS: none
1953  *
1954  * RETURNS:
1955  *   success: FSCK_OK
1956  *   failure: something else
1957  */
1958 int verify_blkall_summary_msgs()
1959 {
1960     int vbsm_rc = FSCK_OK;
1961
1962     if (bmap_recptr->dmap_pmap_error) {
1963         fsck_send_msg(fsck_BADDMAPPMAPS, 0);
1964     }
1965     if (bmap_recptr->dmap_slfv_error) {
1966         msgprms[0] = message_parm_0;
1967         msgprmidx[0] = fsck_dmap;
1968         fsck_send_msg(fsck_BADBMAPSLFV, 1);
1969     }
1970     if (bmap_recptr->dmap_slnv_error) {
1971         msgprms[0] = message_parm_0;
1972         msgprmidx[0] = fsck_dmap;
1973         fsck_send_msg(fsck_BADBMAPSLNV, 1);
1974     }
1975     if (bmap_recptr->dmap_other_error) {
1976         msgprms[0] = message_parm_0;
1977         msgprmidx[0] = fsck_dmap;
1978         fsck_send_msg(fsck_BADBMAPSOTHER, 1);

```

```

1978     }
1979     if (bmap_recptr->L0pg_slfv_error) {
1980         msgprms[0] = message_parm_0;
1981         msgprmidx[0] = fsck_L0;
1982         fsck_send_msg(fsck_BADBMAPSLFV, 1);
1983     }
1984     if (bmap_recptr->L0pg_slnv_error) {
1985         msgprms[0] = message_parm_0;
1986         msgprmidx[0] = fsck_L0;
1987         fsck_send_msg(fsck_BADBMAPSLNV, 1);
1988     }
1989     if (bmap_recptr->L0pg_other_error) {
1990         msgprms[0] = message_parm_0;
1991         msgprmidx[0] = fsck_L0;
1992         fsck_send_msg(fsck_BADBMAPSOTHER, 1);
1993     }
1994     if (bmap_recptr->L1pg_slfv_error) {
1995         msgprms[0] = message_parm_0;
1996         msgprmidx[0] = fsck_L1;
1997         fsck_send_msg(fsck_BADBMAPSLFV, 1);
1998     }
1999     if (bmap_recptr->L1pg_slnv_error) {
2000         msgprms[0] = message_parm_0;
2001         msgprmidx[0] = fsck_L1;
2002         fsck_send_msg(fsck_BADBMAPSLNV, 1);
2003     }
2004     if (bmap_recptr->L1pg_other_error) {
2005         msgprms[0] = message_parm_0;
2006         msgprmidx[0] = fsck_L1;
2007         fsck_send_msg(fsck_BADBMAPSOTHER, 1);
2008     }
2009     if (bmap_recptr->L2pg_slfv_error) {
2010         msgprms[0] = message_parm_0;
2011         msgprmidx[0] = fsck_L2;
2012         fsck_send_msg(fsck_BADBMAPSLFV, 1);
2013     }
2014     if (bmap_recptr->L2pg_slnv_error) {
2015         msgprms[0] = message_parm_0;
2016         msgprmidx[0] = fsck_L2;
2017         fsck_send_msg(fsck_BADBMAPSLNV, 1);
2018     }
2019     if (bmap_recptr->L2pg_other_error) {
2020         msgprms[0] = message_parm_0;
2021         msgprmidx[0] = fsck_L2;
2022         fsck_send_msg(fsck_BADBMAPSOTHER, 1);
2023     }
2024     if (bmap_recptr->ctl_fctl_error) {
2025         fsck_send_msg(fsck_BADBMAPCAGFCL, 0);
2026     }
2027     if (bmap_recptr->ctl_other_error) {
2028         fsck_send_msg(fsck_BADBMAPCOTH, 0);
2029     }
2030     if (bmap_recptr->dmap_pmap_error || bmap_recptr->dmap_slfv_error ||
2031         bmap_recptr->dmap_slnv_error || bmap_recptr->dmap_other_error ||
2032         bmap_recptr->L0pg_slfv_error || bmap_recptr->L0pg_slnv_error ||
2033         bmap_recptr->L0pg_other_error || bmap_recptr->L1pg_slfv_error ||
2034         bmap_recptr->L1pg_slnv_error || bmap_recptr->L1pg_other_error ||
2035         bmap_recptr->L2pg_slfv_error || bmap_recptr->L2pg_slnv_error ||
2036         bmap_recptr->L2pg_other_error) {
2037         agg_recptr->ag_dirty = 1;
2038         fsck_send_msg(fsck_BADBLKALLOC, 0);
2039     }
2040     if (bmap_recptr->ctl_fctl_error || bmap_recptr->ctl_other_error) {
2041         agg_recptr->ag_dirty = 1;
2042         fsck_send_msg(fsck_BADBLKALLOCCTL, 0);
2043     }
2044     return (vbsm_rc);
2045 }

```



```
1 /*
2  * Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  * This program is free software; you can redistribute it and/or modify
5  * it under the terms of the GNU General Public License as published by
6  * the Free Software Foundation; either version 2 of the License, or
7  * (at your option) any later version.
8  *
9  * This program is distributed in the hope that it will be useful,
10 * but WITHOUT ANY WARRANTY; without even the implied warranty of
11 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 * the GNU General Public License for more details.
13 *
14 * You should have received a copy of the GNU General Public License
15 * along with this program; if not, write to the Free Software
16 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #include "xfskint.h"
19
20 /* + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
21  *
22  * superblock buffer pointer
23  *
24  * defined in xchkdsk.c
25  */
26 extern struct superblock *sb_ptr;
27
28 /* + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
29  *
30  * fsck aggregate info structure pointer
31  *
32  * defined in xchkdsk.c
33  */
34 extern struct fsck_agg_record *agg_recptr;
35
36 /* + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
37  *
38  * For message processing
39  *
40  * defined in xchkdsk.c
41  */
42 extern char message_parm_0[];
43 extern char message_parm_1[];
44 extern char message_parm_2[];
45 extern char message_parm_3[];
46 extern char message_parm_4[];
47 extern char message_parm_5[];
48 extern char message_parm_6[];
49 extern char message_parm_7[];
50 extern char message_parm_8[];
51 extern char message_parm_9[];
52
53 extern char *msgprms[];
54 extern int16_t msgprmidx[];
55
56 extern char *verbose_msg_ptr;
57
58 extern char *MsgText[];
59
60 extern char *Vol_Label;
61
62 /* + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
63  *
64  * For directory entry processing
65  *
66  * defined in xchkdsk.c
67  */
68 extern uint32_t key_len[2];
69 extern UniChar key[2][JFS_NAME_MAX];
70 extern UniChar ukey[2][JFS_NAME_MAX];
71
72 extern int32_t Uni_Name_len;
73 extern UniChar Uni_Name[JFS_NAME_MAX];
74
75 /* + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
76  *
77  * Device information.
78  *
79  * defined in xchkdsk.c
80  */
81 extern HFILE Dev_IOPort;
82 extern uint32_t Dev_blksize;
83
84 /* + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + + +
85  *
86  * The current processing phase.
87  *
88  * defined in xchkdsk.c
89  * constants defining possible values are defined in xfsc.h
90  */
```



```

181  *      success: FSCK_OK
182  *      failure: something else
183  */
184  int adjust_parents(struct fsck_inode_record *ino_recptr, uint32_t ino_idx)
185  {
186      int adjps_rc = FSCK_OK;
187      struct fsck_inode_ext_record *this_ext;
188      struct fsck_inode_ext_record *rest_of_list;
189      int keep_primary_parent = 0;
190
191      /*
192       * if this is a directory with illegal hard links then the
193       * inode number in the fsck inode record is the one stored in
194       * the inode on disk. Then if the inode isn't being released
195       * and any observed parent matches the stored parent, that
196       * parent will not be adjusted.
197       */
198      if ((ino_recptr->inode_type == directory_inode)
199          && (ino_recptr->unxpctd_prnts)) {
200          /* dir with multiple parents */
201          if (!ino_recptr->selected_to_rls) {
202              /* not being released */
203              if ((ino_recptr->parent_inonum != ROOT_I)
204                  || (!agg_recptr->rootdir_rebuilt)) {
205                  keep_primary_parent = 1;
206              }
207          }
208      } else {
209          /* not a dir with multiple parents -- this must be either an
210           * inode approved for released or an unallocated inode with
211           * parents observed.
212           */
213          /*
214           * the 1st parent observed is in the inode record. Any others are
215           * in extension records.
216           */
217          if (ino_recptr->parent_inonum != 0) {
218              if ((ino_recptr->parent_inonum != ROOT_I)
219                  || (!agg_recptr->rootdir_rebuilt)) {
220                  /* either this parent isn't the root or
221                   * else the root dir has not been rebuilt */
222                  adjps_rc =
223                      adjust_parent(ino_idx, ino_recptr,
224                                   ino_recptr->parent_inonum);
225                  ino_recptr->parent_inonum = 0; /* clear it */
226              }
227          }
228      }
229
230      /*
231       * detach the extensions list from the inode record
232       */
233      this_ext = ino_recptr->ext_rec;
234      ino_recptr->ext_rec = NULL;
235
236      while ((adjps_rc == FSCK_OK) && (this_ext != NULL)) {
237          /* there may be more parents */
238          rest_of_list = this_ext->next;
239          if (this_ext->ext_type != parent_extension) {
240              /* not a parent */
241              this_ext->next = ino_recptr->ext_rec;
242              ino_recptr->ext_rec = this_ext;
243          } else {
244              /* parent extension */
245              if ((this_ext->inonum == ROOT_I)
246                  && (agg_recptr->rootdir_rebuilt)) {
247                  /*
248                   * This parent is the root and the root dir has
249                   * been rebuilt. This is equivalent to a parent
250                   * marked for release.
251                   */
252                  release_inode_extension(this_ext);
253              } else if ((keep_primary_parent) &&
254                         (this_ext->inonum ==
255                          ino_recptr->parent_inonum)) {
256                  /*
257                   * We're keeping the entry for the expected parent.
258                   * Just drop the extension record and clear the
259                   * 'unexpected parents' flag. When we finish this
260                   * routine the inode will be all set.
261                   */
262                  release_inode_extension(this_ext);
263                  ino_recptr->unxpctd_prnts = 0;
264              } else {
265                  /* either not keeping the 'primary'
266                   * or else this isn't it */
267                  adjps_rc =
268                      adjust_parent(ino_idx, ino_recptr,

```

```

271                 this_ext->inonum);
272                 release_inode_extension(this_ext);
273             }
274         }
275
276         this_ext = rest_of_list;
277     }
278     return (adjps_rc);
279 }
280
281 /*****
282  * NAME: check_connectedness
283  *
284  * FUNCTION: Verify that, after approved corrections are made, all inodes
285  *           in use will be connected to the root directory tree.
286  *
287  * PARAMETERS: none
288  *
289  * NOTES:    o A directory inode must have exactly one parent inode.
290  *           o A non-directory inode must have at least one parent inode.
291  *
292  * RETURNS:
293  *           success: FSCK_OK
294  *           failure: something else
295  */
296 int check_connectedness()
297 {
298     int cc_rc = FSCK_OK;
299     uint32_t ino_idx;
300     struct fsck_inode_record *this_inorec;
301     struct fsck_inode_ext_record *new_ext;
302     int aggregate_inode = 0;
303
304     struct fsck_ino_msg_info ino_msg_info;
305     struct fsck_ino_msg_info *msg_info_ptr;
306
307     msg_info_ptr = &ino_msg_info;
308     /* all fileset owned */
309     msg_info_ptr->msg_inopfx = fsck_fset_inode;
310
311     /*
312      * detect orphan inodes, including ones which we're about to
313      * orphan by releasing inodes.
314      *
315      * if any non-orphan inode was flagged as a directory with illegal
316      * hard links, see if it's going to be true after we release inodes.
317      * If not, make sure the remaining link matches the one stored in
318      * the inode.
319      */
320     cc_rc = get_inorecptr_first(aggregate_inode, &ino_idx, &this_inorec);
321
322     while ((cc_rc == FSCK_OK) && (this_inorec != NULL)
323           && (ino_idx < FILESET_OBJECT_I)) {
324         /*
325          * not interesting until we get past the root inode
326          * and the special fileset inodes.
327          */
328         cc_rc =
329             get_inorecptr_next(aggregate_inode, &ino_idx, &this_inorec);
330     }
331
332     while ((cc_rc == FSCK_OK) && (this_inorec != NULL)) {
333
334         if ((this_inorec->in_use) && (!this_inorec->selected_to_rls)) {
335             /* inode in use and not selected to release */
336             msg_info_ptr->msg_inonum = ino_idx;
337             if (this_inorec->inode_type == directory_inode) {
338                 msg_info_ptr->msg_inotyp = fsck_directory;
339
340             } else if (this_inorec->inode_type == symlink_inode) {
341                 msg_info_ptr->msg_inotyp = fsck_symbolic_link;
342
343             } else if (this_inorec->inode_type ==
344                       char_special_inode) {
345                 msg_info_ptr->msg_inotyp = fsck_char_special;
346
347             } else if (this_inorec->inode_type ==
348                       block_special_inode) {
349                 msg_info_ptr->msg_inotyp = fsck_block_special;
350
351             } else if (this_inorec->inode_type == FIFO_inode) {
352                 msg_info_ptr->msg_inotyp = fsck_FIFO;
353
354             } else if (this_inorec->inode_type == SOCK_inode) {
355                 msg_info_ptr->msg_inotyp = fsck SOCK;
356
357             } else {
358                 /* a regular file */
359                 msg_info_ptr->msg_inotyp = fsck_file;
360             }
361         }
362     }

```

```

361
362         if (this_inorec->parent_inonum == 0) {
363             /* no parents were observed by fsck */
364             msgprms[0] = message_parm_0;
365             msgprmidx[0] = msg_info_ptr->msg_inopfx;
366             sprintf(message_parm_1, "%d",
367                 msg_info_ptr->msg_inonum);
368             msgprms[1] = message_parm_1;
369             msgprmidx[1] = 0;
370             fsck_send_msg(fsck_INONOPATHS, 2);
371         } else {
372             /* at least one parent observed by fsck */
373             /*
374              * make adjustments to child records for
375              * parents which will be released
376              */
377             cc_rc = reset_parents(this_inorec, ino_idx);
378         }
379
380     if (cc_rc == FSCK_OK) {
381         if (this_inorec->parent_inonum == 0) {
382             /* unconnected !! */
383             /*
384              * do not issue a message to inform the user about
385              * this condition since it is a side effect of
386              * the (approved) release of some other inode(s)
387              */
388             this_inorec->reconnect = 1;
389             agg_recptr->corrections_approved = 1;
390             cc_rc = get_inode_extension(&new_ext);
391             if (cc_rc == FSCK_OK) {
392                 /* got an extension record */
393                 new_ext->ext_type =
394                     add_direntry_extension;
395                 new_ext->inonum = ino_idx;
396                 new_ext->ino_type =
397                     this_inorec->inode_type;
398                 new_ext->next =
399                     agg_recptr->
400                     inode_reconn_extens;
401                 agg_recptr->
402                     inode_reconn_extens =
403                     new_ext;
404                 /* increment for the link from
405                  * parent after reconnect
406                  */
407                 this_inorec->link_count++;
408             }
409         } else {
410             /* else still connected */
411             if (this_inorec->unxpctd_prnts) {
412                 /* multiple parents */
413                 cc_rc =
414                     display_paths(ino_idx,
415                                 this_inorec,
416                                 msg_info_ptr);
417
418                 if (cc_rc == FSCK_OK) {
419
420                     if (agg_recptr->
421                         processing_readwrite)
422                     {
423                         cc_rc =
424                             adjust_parents
425                                 (this_inorec,
426                                 ino_idx);
427                         msgprms[0] =
428                             message_parm_0;
429                         msgprmidx[0] =
430                             msg_info_ptr->
431                             msg_inopfx;
432                         sprintf
433                             (message_parm_1,
434                              "%d",
435                               msg_info_ptr->
436                               msg_inonum);
437                         msgprms[1] =
438                             message_parm_1;
439                         msgprmidx[1] =
440                             0;
441                         fsck_send_msg
442                             (fsck_WILLFIXDIRWHDLKS,
443                              2);
444                     } else {
445                         /* no write access */
446                         this_inorec->
447                             unxpctd_prnts
448                             = 0;
449                         agg_recptr->

```

```

451         ag_dirty =
452         1;
453         msgprms[0] =
454         message_parm_0;
455         msgprmidx[0] =
456         msg_info_ptr->
457         msg_inopfx;
458         sprintf
459         (message_parm_1,
460         "%d",
461         msg_info_ptr->
462         msg_inonum);
463         msgprms[1] =
464         message_parm_1;
465         msgprmidx[1] =
466         0;
467         fsck_send_msg
468         (fsck_DIRWHDLKS,
469         2);
470     }
471 }
472 } else if ((cc_rc == FSCK_OK)
473           && (this_inorec->
474             crrct_prnt_inonum)) {
475     /*
476     * a single parent but not the one
477     * named in the implied '..' entry
478     */
479     cc_rc =
480     display_paths(ino_idx,
481                 this_inorec,
482                 msg_info_ptr);
483
484     if (agg_recptr->
485         processing_readwrite) {
486         msgprms[0] =
487         message_parm_0;
488         msgprmidx[0] =
489         msg_info_ptr->
490         msg_inopfx;
491         sprintf(message_parm_1,
492               "%d",
493               msg_info_ptr->
494               msg_inonum);
495         msgprms[1] =
496         message_parm_1;
497         msgprmidx[1] = 0;
498         fsck_send_msg
499         (fsck_WILLFIXINCREF,
500         2);
501     } else {
502         /* no write access */
503         this_inorec->
504         crrct_prnt_inonum =
505         0;
506         agg_recptr->ag_dirty =
507         1;
508         msgprms[0] =
509         message_parm_0;
510         msgprmidx[0] =
511         msg_info_ptr->
512         msg_inopfx;
513         sprintf(message_parm_1,
514               "%d",
515               msg_info_ptr->
516               msg_inonum);
517         msgprms[1] =
518         message_parm_1;
519         msgprmidx[1] = 0;
520         fsck_send_msg
521         (fsck_INCIgnoref, 2);
522     }
523 }
524 }
525 }
526 }
527 if (cc_rc == FSCK_OK) {
528     cc_rc =
529     get_inorecptr_next(aggregate_inode, &ino_idx,
530                       &this_inorec);
531 }
532 }
533 return (cc_rc);
534 }
535
536 /*****
537 * NAME: check_dir_integrity
538 *
539 * FUNCTION: Verify that no directory has more than 1 entry for any
540 *           single inode. If a directory does, then that directory

```

```

541 *           is corrupt.
542 *
543 * RETURNS:
544 *   success: FSCK_OK
545 *   failure: something else
546 */
547 int check_dir_integrity()
548 {
549     int cdi_rc = FSCK_OK;
550     uint32_t ino_idx;
551     struct fsck_inode_record *this_inorec;
552     struct fsck_inode_record *that_inorec;
553     int aggregate_inode = 0;
554     int alloc_ifnull = 0;
555     struct fsck_inode_ext_record *this_ext;
556     struct fsck_inode_ext_record *that_ext;
557     int dup_parent_detected = 0;
558
559     struct fsck_ino_msg_info ino_msg_info;
560     struct fsck_ino_msg_info *msg_info_ptr;
561
562     msg_info_ptr = &ino_msg_info;
563     /* all fileset owned */
564     msg_info_ptr->msg_inopfx = fsck_fset_inode;
565
566     /*
567      * Verify that no inode has multiple links from the same
568      * directory.
569      */
570     cdi_rc = get_inorecptr_first(aggregate_inode, &ino_idx, &this_inorec);
571
572     while ((cdi_rc == FSCK_OK) && (this_inorec != NULL)
573           && (ino_idx < FILESET_OBJECT_I)) {
574         /*
575          * not interesting until we get past the root inode
576          * and the special fileset inodes.
577          */
578         cdi_rc =
579             get_inorecptr_next(aggregate_inode, &ino_idx, &this_inorec);
580     }
581
582     while ((cdi_rc == FSCK_OK) && (this_inorec != NULL)) {
583         if ((this_inorec->in_use) && (!this_inorec->selected_to_rls)
584             && (this_inorec->inode_type == directory_inode)) {
585             /* directory inode in use and not selected to release */
586             msg_info_ptr->msg_inotyp = fsck_directory;
587             msg_info_ptr->msg_inonum = ino_idx;
588
589             if (this_inorec->parent_inonum != 0) {
590                 /* at least 1 parent observed by fsck */
591                 if (this_inorec->ext_rec) {
592                     /* maybe more parents have been observed */
593                     /*
594                      * get the first entry in the extensions
595                      * list on the inode record
596                      */
597                     this_ext = this_inorec->ext_rec;
598
599                     while ((cdi_rc == FSCK_OK)
600                           && (this_ext != NULL)) {
601                         /* there may be more parents */
602                         if (this_ext->ext_type ==
603                             parent_extension) {
604                             /* a parent */
605                             cdi_rc =
606                                 get_inorecptr
607                                 (aggregate_inode,
608                                 alloc_ifnull,
609                                 this_ext->inonum,
610                                 &that_inorec);
611                             if ((cdi_rc == FSCK_OK)
612                                 && (!that_inorec->
613                                     selected_to_rls))
614                                 {
615                                     /* parent isn't marked for release (yet) */
616                                     dup_parent_detected
617                                         = 0;
618                                     if (this_ext->
619                                         inonum ==
620                                         this_inorec->
621                                         parent_inonum)
622                                     {
623                                         dup_parent_detected
624                                             =
625                                             -1;
626                                     } else {
627                                         /* need to check for dups in rest of list
628
629                                         that_ext
630                                             =

```

```

630         this_ext->
631         next;
632         while ((!dup_parent_detected) && (that_ext
633             if (that_ext->ext_type == parent_
634                 /* another parent extensi
635                 if (this_ext->inonum == t
636                     dup_parent_detect
637                         =
638                         -1;
639             }
640             that_ext
641             =
642             that_ext->
643             next;
644         }
645     }
646     if (dup_parent_detected) {
647         /*
648         * mark the parent's inode record for rel
649         ease
650         */
651         that_inorec->
652         selected_to_rls
653         = 1;
654         /*
655         * notify the user that the directory is
656         bad
657         */
658         msgprms
659         [0]
660         =
661         message_parm_0;
662         msgprmidx
663         [0]
664         =
665         fsck_directory;
666         msgprms
667         [1]
668         =
669         message_parm_1;
670         msgprmidx
671         [1]
672         =
673         msg_info_ptr->
674         msg_inopfx;
675         sprintf
676         (message_parm_2,
677          "%d",
678          this_ext->
679          inonum);
680         msgprms
681         [2]
682         =
683         message_parm_2;
684         msgprmidx
685         [2]
686         = 0;
687         sprintf
688         (message_parm_3,
689          "%d",
690          37);
691         msgprms
692         [3]
693         =
694         message_parm_3;
695         msgprmidx
696         [3]
697         = 0;
698         fsck_send_msg
699         (fsck_BADKEYS,
700          4);
701     }
702     }
703     }
704     }
705     }
706     }
707     if (cdi_rc == FSCK_OK) {
708         cdi_rc =
709         get_inorecptr_next(aggregate_inode, &ino_idx,
710                          &this_inorec);
711     }
712 }

```



```

713         return (cdi_rc);
714     }
715 }
716
717 /*****
718  * NAME: check_link_counts
719  *
720  * FUNCTION: Count links from child directories to their parents.
721  *
722  *           Verify that the link count stored in each in-use inode
723  *           matches the number of links fsck observed for the inode.
724  *
725  * RETURNS:
726  *         success: FSCK_OK
727  *         failure: something else
728  */
729 int check_link_counts()
730 {
731     int clc_rc = FSCK_OK;
732     uint32_t ino_idx;
733     int num_parents;
734     int invalid_count_seen = 0;
735     int low_stored_count_seen = 0;
736     struct fsck_inode_ext_record *this_ext;
737     struct fsck_inode_record *this_inorec;
738     struct fsck_inode_record *parent_inorec;
739     int done_looking = 0;
740     int aggregate_inode = 0;
741     int alloc_ifnull = 0;
742
743     struct fsck_ino_msg_info ino_msg_info;
744     struct fsck_ino_msg_info *msg_info_ptr;
745
746     msg_info_ptr = &ino_msg_info;
747     /* all fileset owned */
748     msg_info_ptr->msg_inopfx = fsck_fset_inode;
749
750     /*
751      * count links from child directories to their parents
752      *
753      * (These can't be counted when the parent-child relationship
754      * is observed because that observation occurs while processing
755      * the parent and until the child is processed we don't know
756      * whether the child is a directory or not.)
757      */
758     clc_rc = get_inorecptr_first(aggregate_inode, &ino_idx, &this_inorec);
759
760     while ((clc_rc == FSCK_OK) && (this_inorec != NULL)) {
761         if ((this_inorec->in_use) && (!this_inorec->selected_to_rls)
762             && (!this_inorec->ignore_alloc_blks)
763             && (this_inorec->inode_type == DIRECTORY_INODE)) {
764             /* inode is in use, not being released, and is type directory */
765             /* for the self entry */
766             this_inorec->link_count++;
767             if ((this_inorec->parent_inonum == ROOT_I)
768                 && (agg_recptr->rootdir_rebuilt)) {
769                 /*
770                  * special case: if the parent is root and root was
771                  * rebuilt, then don't increment parent
772                  */
773                 if (this_inorec->inonum == ROOT_I) {
774                     /*
775                      * special case: if this IS the root, then it's
776                      * link from itself to itself DOES count
777                      */
778                     this_inorec->link_count++;
779                 }
780             } else if (this_inorec->parent_inonum != 0) {
781                 /* not an orphan */
782                 clc_rc =
783                     get_inorecptr(aggregate_inode, alloc_ifnull,
784                                 this_inorec->parent_inonum,
785                                 &parent_inorec);
786                 if ((clc_rc != FSCK_OK)
787                     || (parent_inorec == NULL)) {
788                     clc_rc = FSCK_INTERNAL_ERROR_13;
789                     sprintf(message_parm_0, "%d", clc_rc);
790                     msgprms[0] = message_parm_0;
791                     msgprmidx[0] = 0;
792                     sprintf(message_parm_1, "%d", ino_idx);
793                     msgprms[1] = message_parm_1;
794                     msgprmidx[1] = 0;
795                     sprintf(message_parm_2, "%d",
796                             this_inorec->parent_inonum);
797                     msgprms[2] = message_parm_2;
798                     msgprmidx[2] = 0;
799                     sprintf(message_parm_3, "%d", 0);
800                     msgprms[3] = message_parm_3;
801                     msgprmidx[3] = 0;
802                     fsck_send_msg(fsck_INTERNALERROR, 4);

```

```

803     } else {
804         /* handle the first (and usually the only) parent. */
805         parent_inorec->link_count++;
806     }
807     if ((clc_rc == FSCK_OK)
808         && (this_inorec->ext_rec != NULL)) {
809         /* there might be more parents */
810         num_parents = parent_count(this_inorec);
811         if (num_parents > 1) {
812             /* directory with illegal links */
813             this_inorec->unxpctd_prnts = 1;
814             agg_recptr->corrections_needed =
815                 1;
816             /*
817              * Create an extension record for the parent inode
818              * number now stored in the child inode record.
819              * When we traverse the aggregate on-disk we'll copy
820              * the stored value into this field of the inode record
821              * for use when displaying paths to the inode.
822              */
823             clc_rc =
824                 get_inode_extension
825                 (&this_ext);
826             if (clc_rc == FSCK_OK) {
827                 /* got extension record */
828                 this_ext->ext_type =
829                     parent_extension;
830                 this_ext->inonum =
831                     this_inorec->
832                     parent_inonum;
833                 this_ext->next =
834                     this_inorec->
835                     ext_rec;
836                 this_inorec->ext_rec =
837                     this_ext;
838                 this_inorec->
839                     parent_inonum = 0;
840                 /* already counted the first
841                  * one, back when it was in the
842                  * workspace inode record itself
843                  */
844                 this_ext =
845                     this_ext->next;
846                 while ((clc_rc ==
847                        FSCK_OK)
848                        && (this_ext !=
849                           NULL)) {
850                     /* exten records to check */
851                     if (this_ext->
852                         ext_type ==
853                         parent_extension)
854                     {
855                         clc_rc =
856                             get_inorecptr
857                             (aggregate_inode,
858                              alloc_ifnull,
859                              this_ext->
860                              inonum,
861                              &parent_inorec);
862                         if ((clc_rc != FSCK_OK)
863                             ||
864                             (parent_inorec
865                              ==
866                               NULL))
867                         {
868                             clc_rc
869                                 =
870                                 FSCK_INTERNAL_ERROR_14;
871                             sprintf
872                                 (message_parm_0,
873                                  "%d",
874                                  clc_rc);
875                             msgprms
876                                 [0]
877                                 =
878                                 message_parm_0;
879                             msgprmidx
880                                 [0]
881                                 =
882                                 0;
883                             sprintf
884                                 (message_parm_1,
885                                  "%d",
886                                  ino_idx);
887                             msgprms
888                                 [1]
889                                 =
890                                 message_parm_1;
891                             msgprmidx
892                                 [1]

```

```

893                                     =
894                                     0;
895                                     sprintf
896                                     (message_parm_2,
897                                     "%d",
898                                     this_ext->
899                                     inonum);
900                                     msgprms
901                                     [2]
902                                     =
903                                     message_parm_2;
904                                     msgprmidx
905                                     [2]
906                                     =
907                                     0;
908                                     sprintf
909                                     (message_parm_3,
910                                     "%d",
911                                     0);
912                                     msgprms
913                                     [3]
914                                     =
915                                     message_parm_3;
916                                     msgprmidx
917                                     [3]
918                                     =
919                                     0;
920                                     fsck_send_msg
921                                     (fsck_INTERNALERROR,
922                                     4);
923                                     } else {
924                                     parent_inorec->
925                                     link_count++;
926                                     }
927                                     }
928                                     this_ext =
929                                     this_ext->
930                                     next;
931                                     }
932                                     }
933                                     }
934                                     }
935                                     }
936                                     }
937                                     if (clc_rc == FSCK_OK) {
938                                     clc_rc =
939                                     get_inorecptr_next(aggregate_inode, &ino_idx,
940                                     &this_inorec);
941                                     }
942                                     }
943
944                                     /*
945                                     * verify that stored link counts match observed link counts
946                                     *
947                                     * We have added each observed link and subtracted the stored
948                                     * count. If the stored count is correct the result is 0.
949                                     */
950                                     if (clc_rc == FSCK_OK) {
951                                     clc_rc =
952                                     get_inorecptr_first(aggregate_inode, &ino_idx,
953                                     &this_inorec);
954
955                                     while ((clc_rc == FSCK_OK) && (this_inorec != NULL)
956                                     && (!done_looking)) {
957
958                                     if ((this_inorec->in_use)
959                                     && (!this_inorec->selected_to_rls)
960                                     && (!this_inorec->ignore_alloc_blks)) {
961                                     /* inode is in use and not being released */
962                                     if (this_inorec->link_count != 0) {
963                                     /* stored link count doesn't match fsck's observations */
964                                     if (this_inorec->parent_inonum == 0) {
965                                     /* inode is an orphan */
966                                     this_inorec->crrct_link_count =
967                                     1;
968                                     } else {
969                                     /* not an orphan */
970                                     this_inorec->crrct_link_count =
971                                     1;
972                                     if (this_inorec->link_count > 0) {
973                                     low_stored_count_seen =
974                                     1;
975                                     }
976                                     invalid_count_seen = 1;
977
978                                     msgprms[0] = message_parm_0;
979                                     msgprmidx[0] =
980                                     msg_info_ptr->msg_inopfx;
981                                     sprintf(message_parm_1, "%d",
982                                     ino_idx);

```

```

983         msgprms[1] = message_parm_1;
984         msgprmidx[1] = 0;
985         fsck_send_msg(fsck_BADINOLKCT,
986                     2);
987     }
988 }
989
990     if (clc_rc == FSCK_OK) {
991         clc_rc =
992             get_inorecptr_next(aggregate_inode,
993                             &ino_idx, &this_inorec);
994     }
995 }
996
997     if ((clc_rc == FSCK_OK) && (invalid_count_seen)) {
998
999         if (agg_recptr->processing_readwrite) {
1000             agg_recptr->corrections_approved = 1;
1001             fsck_send_msg(fsck_WILLFIXLINKCTS, 0);
1002         } else {
1003             /* no write access */
1004             if (low_stored_count_seen) {
1005                 agg_recptr->ag_dirty = 1;
1006             }
1007             /*
1008              * reset all link counts (in the fsck workspace) to
1009              * zero so that we won't accidentally correct them
1010              * while doing link count adjustments.
1011              *
1012              * (Link count adjustments are side effects of approved
1013              * repairs. For example, if a directory inode is
1014              * released, the link count of its parent directory
1015              * is decremented.)
1016              */
1017             clc_rc =
1018                 get_inorecptr_first(aggregate_inode,
1019                                   &ino_idx, &this_inorec);
1020
1021             while ((clc_rc == FSCK_OK)
1022                  && (this_inorec != NULL)) {
1023
1024                 if (this_inorec->in_use) {
1025                     this_inorec->corrct_link_count =
1026                         0;
1027                     this_inorec->link_count = 0;
1028                 }
1029                 clc_rc =
1030                     get_inorecptr_next(aggregate_inode,
1031                                       &ino_idx,
1032                                       &this_inorec);
1033             }
1034             fsck_send_msg(fsck_BADLINKCTS, 0);
1035         }
1036     }
1037 }
1038     return (clc_rc);
1039 }
1040
1041 /*****
1042  * NAME: reset_parents
1043  *
1044  * FUNCTION: Adjust the fsck notations about the inode's parent(s) if
1045  *           the parent(s) are corrupt or approved for release.
1046  *
1047  * PARAMETERS:
1048  *     ino_recptr - input - pointer to an fsck record describing the inode
1049  *     ino_idx    - input - ordinal number of the inode
1050  *
1051  * RETURNS:
1052  *     success: FSCK_OK
1053  *     failure: something else
1054  */
1055 int reset_parents(struct fsck_inode_record *ino_recptr, uint32_t ino_idx)
1056 {
1057     int resps_rc = FSCK_OK;
1058     struct fsck_inode_ext_record *this_ext;
1059     struct fsck_inode_ext_record *rest_of_list;
1060     int parent_count = 0;
1061     uint32_t stored_parent_inonum = 0;
1062     struct fsck_inode_record *parent_inorecptr;
1063     int aggregate_inode = 0;
1064     int alloc_ifnull = 0;
1065     /*
1066      * if this is a directory with illegal hard links the inode
1067      * number in the fsck inode record is the one stored in the
1068      * inode on disk.
1069      */
1070     if ((ino_recptr->inode_type == directory_inode)
1071         && (ino_recptr->unxpctd_prnts)) {
1072         /* dir with multiple parents */

```

```

1073         /*
1074         * Save the value stored in the inode record and then clear it.
1075         */
1076         stored_parent_inonum = ino_recptr->parent_inonum;
1077         ino_recptr->parent_inonum = 0;
1078     } else {
1079         /* not a dir with multiple parents */
1080         /*
1081         * the 1st parent observed is in the inode record.
1082         * Any others are in extension records.
1083         */
1084         resps_rc = get_inorecptr(aggregate_inode, alloc_ifnull,
1085                                 ino_recptr->parent_inonum,
1086                                 &parent_inorecptr);
1087         if ((resps_rc == FSCK_OK) && (parent_inorecptr == NULL)) {
1088             resps_rc = FSCK_INTERNAL_ERROR_15;
1089             sprintf(message_parm_0, "%d", resps_rc);
1090             msgprms[0] = message_parm_0;
1091             msgprmidx[0] = 0;
1092             sprintf(message_parm_1, "%d", ino_idx);
1093             msgprms[1] = message_parm_1;
1094             msgprmidx[1] = 0;
1095             sprintf(message_parm_2, "%d",
1096                     ino_recptr->parent_inonum);
1097             msgprms[2] = message_parm_2;
1098             msgprmidx[2] = 0;
1099             sprintf(message_parm_3, "%d", 0);
1100             msgprms[3] = message_parm_3;
1101             msgprmidx[3] = 0;
1102             fsck_send_msg(fsck_INTERNALERROR, 4);
1103         }
1104         else if (resps_rc == FSCK_OK) {
1105             if ((ino_recptr->parent_inonum == ROOT_I)
1106                 && (agg_recptr->rootdir_rebuilt)) {
1107                 /*
1108                 * special case: if the parent is root and root
1109                 * was rebuilt, then this is an orphan
1110                 */
1111                 ino_recptr->parent_inonum = 0;
1112                 ino_recptr->link_count--;
1113             } else if ((!parent_inorecptr->selected_to_rls)
1114                       && (!parent_inorecptr->ignore_alloc_blks)) {
1115                 /*
1116                 * keeping this parent and haven't found the
1117                 * tree to be corrupt
1118                 */
1119                 parent_count++;
1120             }
1121             else {
1122                 /* releasing this parent */
1123                 ino_recptr->parent_inonum = 0;
1124                 ino_recptr->link_count--;
1125             }
1126         }
1127     }
1128     /*
1129     * detach the extensions list from the inode record
1130     */
1131     this_ext = ino_recptr->ext_rec;
1132     ino_recptr->ext_rec = NULL;
1133     while ((resps_rc == FSCK_OK) && (this_ext != NULL)) {
1134         /* there may be more parents */
1135         rest_of_list = this_ext->next;
1136         if (this_ext->ext_type != parent_extension) {
1137             /* not a parent */
1138             this_ext->next = ino_recptr->ext_rec;
1139             ino_recptr->ext_rec = this_ext;
1140         } else {
1141             /* parent extension */
1142             resps_rc = get_inorecptr(aggregate_inode, alloc_ifnull,
1143                                     this_ext->inonum,
1144                                     &parent_inorecptr);
1145             if ((resps_rc == FSCK_OK) && (parent_inorecptr == NULL)) {
1146                 resps_rc = FSCK_INTERNAL_ERROR_16;
1147                 sprintf(message_parm_0, "%d", resps_rc);
1148                 msgprms[0] = message_parm_0;
1149                 msgprmidx[0] = 0;
1150                 sprintf(message_parm_1, "%d", ino_idx);
1151                 msgprms[1] = message_parm_1;
1152                 msgprmidx[1] = 0;
1153                 sprintf(message_parm_2, "%d", this_ext->inonum);
1154                 msgprms[2] = message_parm_2;
1155                 msgprmidx[2] = 0;
1156                 sprintf(message_parm_3, "%d", 0);
1157                 msgprms[3] = message_parm_3;
1158                 msgprmidx[3] = 0;
1159                 fsck_send_msg(fsck_INTERNALERROR, 4);
1160             }
1161             else if (resps_rc == FSCK_OK) {
1162                 if ((ino_recptr->parent_inonum == ROOT_I)

```

```

1163         && (agg_recptr->rootdir_rebuilt)) {
1164             /*
1165              * special case: if the parent is root and
1166              * root was rebuilt, then this is an orphan
1167              */
1168             release_inode_extension(this_ext);
1169             ino_recptr->link_count--;
1170
1171         } else if ((!parent_inorecptr->selected_to_rls)
1172                 && (!parent_inorecptr->
1173                    ignore_alloc_blks)) {
1174             /* keeping this parent */
1175             parent_count++;
1176             if (ino_recptr->parent_inonum == 0) {
1177                 ino_recptr->parent_inonum =
1178                     this_ext->inonum;
1179                 release_inode_extension
1180                     (this_ext);
1181             } else {
1182                 /* put it back on the list */
1183                 this_ext->next =
1184                     ino_recptr->ext_rec;
1185                 ino_recptr->ext_rec = this_ext;
1186             }
1187         } else {
1188             /* releasing this parent */
1189             release_inode_extension(this_ext);
1190             ino_recptr->link_count--;
1191         }
1192     }
1193 }
1194 this_ext = rest_of_list;
1195 }
1196 /*
1197 * at this point, if there is at least 1 observed parent which
1198 * is not being released, then a parent inode number is stored in
1199 * the inode record and any other parents are described in extension
1200 * records.
1201 *
1202 * if this is not a directory inode, we're done.
1203 *
1204 * if this is a directory inode, need to recheck for illegal hard
1205 * links and incorrect parent inode entry.
1206 */
1207
1208 if ((resps_rc == FSCK_OK)
1209     && (ino_recptr->inode_type == directory_inode)) {
1210     /* a directory */
1211     if (parent_count == 1) {
1212         /* 1 parent now */
1213         if (ino_recptr->unxpctd_prnts) {
1214             /* entered with multiple links */
1215             /* reset flag */
1216             ino_recptr->unxpctd_prnts = 0;
1217             if (ino_recptr->parent_inonum !=
1218                 stored_parent_inonum) {
1219                 /*
1220                  * Remaining parent doesn't match the one
1221                  * the on-device inode says owns it.
1222                  */
1223                 ino_recptr->corrct_prnt_inonum = 1;
1224                 agg_recptr->corrections_needed = 1;
1225                 agg_recptr->corrections_approved = 1;
1226             }
1227         }
1228     } else if (parent_count == 0) {
1229         /* no parents now */
1230         ino_recptr->corrct_prnt_inonum = 0;
1231         ino_recptr->unxpctd_prnts = 0;
1232     } else {
1233         /* multiple parents still */
1234         ino_recptr->unxpctd_prnts = 1;
1235         agg_recptr->corrections_needed = 1;
1236         resps_rc = get_inode_extension(&this_ext);
1237         if (resps_rc == FSCK_OK) {
1238             this_ext->ext_type = parent_extension;
1239             this_ext->inonum = ino_recptr->parent_inonum;
1240             this_ext->next = ino_recptr->ext_rec;
1241             ino_recptr->ext_rec = this_ext;
1242             ino_recptr->parent_inonum =
1243                 stored_parent_inonum;
1244         }
1245     }
1246 }
1247 return (resps_rc);
1248 }

```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_FSCKPFS
19 #define H_FSCKPFS
20
21 #define fsck_READ 1
22 #define fsck_WRITE 2
23
24 #endif
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_PUFS_CHKDSK
19 #define H_PUFS_CHKDSK
20
21 #include "xfscck.h"
22
23 enum xchkdsk_options {
24     UFS_CHKDSK_LEVEL0,
25     UFS_CHKDSK_LEVEL1,
26     UFS_CHKDSK_LEVEL2,
27     UFS_CHKDSK_LEVEL3,
28     UFS_CHKDSK_IFDIRTY,
29     UFS_CHKDSK_VERBOSE,
30     UFS_CHKDSK_DEBUG,
31     UFS_CHKDSK_CLRBDLKLST,
32     UFS_CHKDSK_SKIPLOGREDO,
33     UFS_CHKDSK_BLOCK_NUMBER,
34     UFS_CHKDSK_INODE_NUMBER,
35     UFS_CHKDSK_FILENAME,
36     UFS_CHKDSK_OPTIONS
37 };
38
39 #endif
```



```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_XFSCCK
19 #define H_XFSCCK
20
21 /*
22  * ----- system includes -----
23  */
24
25 #include <stdarg.h>
26 #include <stddef.h>
27
28 #ifndef _JFS_FSCKDIRE
29 #include <stdio.h>
30 #endif
31
32 #include <string.h>
33 #include <stdlib.h>
34
35 /*
36  * ----- JFS includes -----
37  */
38
39 #include "jfs_types.h"
40 #include "jfs_dinode.h"
41 #include "jfs_dmap.h"
42 #include "jfs_filsys.h"
43 #include "jfs_imap.h"
44 #include "jfs_superblock.h"
45 #include "fsck_base.h"
46 #include "fsckcbb.h"
47
48 /*
49  * jfs library routine includes
50  */
51 #include "libjufs.h"
52 #include "devices.h"
53 #include "diskmap.h"
54 #include "inode.h"
55 #include "logform.h"
56 #include "logredo.h"
57 #include "message.h"
58 #include "super.h"
59
60 /* Stuff for handling extended attributes. */
61 #ifndef OS2
62
63 struct FEA {
64     uint8_t fea;           /* fea */
65     uint8_t cbName;       /* flags */
66     uint16_t cbValue;     /* name length not including NULL */
67 };
68
69 /* flags for FEA.fea */
70 #define FEA_NEEDEA        0x80 /* need EA bit */
71
72 struct FEALIST {
73     uint32_t cbList;     /* feal */
74     struct FEA list[1]; /* total bytes of structure including full list */
75 };
76
77 #define ERROR_EA_LIST_INCONSISTENT 255
78
79 #endif
80
81 extern int jfs_ValidateFEAList(struct FEALIST *, unsigned long *);
82
83 /* ***** IMPORTANT ***** IMPORTANT ***** IMPORTANT ***** IMPORTANT *****
84  *
85  * fsck_first_msgid
86  *   MUST be set to the first message id for fsck, according
87  *   to jfs.txt. Will be used to locate the text for any
88  *   message which is displayed in the local language.
89  *
90  * fsck_highest_msgid_defined

```

```

91  *      MUST be maintained in synch with the
92  *      message id constants (defined in fsckmsgc.h) since the
93  *      message text array and the message attributes array (declared
94  *      in fsckmsgp.h) are both dimensioned using it.
95  *
96  * ***** IMPORTANT ***** IMPORTANT ***** IMPORTANT ***** IMPORTANT ***** */
97
98 #define fsck_msgid_offset          50
99 #define fsck_highest_msgid_defined 599
100
101 #define MAXPARMLEN 64
102
103 #define fscklog_var_text 1
104 #define fscklog_literal 2
105
106 /* + + + + +
107  *
108  * The following are used to access the columns of
109  * MsgProtocol[[]], which is defined in fsckmsgp.h
110  *
111  */
112 #define MP_MSGLVL      0
113 #define MP_MSGFILE     1
114
115 /* + + + + +
116  *
117  * The following are the possible values effective fsck
118  * messaging level and for implied response level.
119  *
120  */
121
122 /*
123  * special cases (must not match any other constants in the table)
124  */
125 #define fsck_hrtbt      130
126 #define fsck_txtins    131
127
128 /*
129  * The lowest messaging level (dictated by input parms) at which
130  * the message is displayed.
131  */
132 #define fsck_quiet      2
133 #define fsck_verbose    8
134 #define fsck_debug     32
135
136 /*
137  * The message file to use for local language lookup
138  *
139  * These are determined by the array message_file_name[]
140  * in messages.c
141  */
142 #define no_msgfile      -1
143 #define jfs_msgfile     1
144
145 /* + + + + +
146  * The following are used for reporting storage allocation
147  * failures.
148  */
149
150 #ifndef _JFS_FSCKDIRE
151 void report_dynstg_error(void);
152
153 #define dynstg_unknown      0
154
155 /* the actions */
156
157 #define dynstg_allocation      1
158 #define dynstg_initialization  2
159
160 /* the objects */
161
162 #define dynstg_unknown_object  0
163 #define dynstg_blkmap          1
164 #define dynstg_blkmap_buf     2
165 #define dynstg_blkmap_hdr     3
166 #define dynstg_inomap         4
167 #define dynstg_fer            5
168 #define dynstg_wspxt          6
169 #define dynstg_pathbuffer     7
170 #define dynstg_inoextrec      8
171 #define dynstg_inorec         9
172 #define dynstg_dtreeQ_elem   10
173 #define dynstg_treeQ_elem    11
174 #define dynstg_ait_map       12
175 #define dynstg_fsit_map      13
176 #define dynstg_dupall_blkrec 14
177 #define dynstg_dupall_inorec 15
178 #define dynstg_agg_agtbl     16
179 #define dynstg_agg_iagtbl    17
180

```



```

361 #define FSCK_ERRORS_UNCORRECTED 4
362 #define FSCK_OP_ERROR 8
363 #define FSCK_USAGE_ERROR 16
364
365 /* informational return codes */
366
367 #define FSCK_FAILED -00001
368
369 #define FSCK_AGFS_INOBAD 10001
370 #define FSCK_AGGAITINOBAD 10002
371 #define FSCK_AGGFSINOBAD 10003
372 #define FSCK_BADAGFSSIZ 10004
373 #define FSCK_BADBLKCTTTL 10005
374 #define FSCK_BADEADESCRIPTOR 10006
375 #define FSCK_BADMDDATA 10007
376 #define FSCK_BADMDDATAEXT 10008
377 #define FSCK_BADMDDATAIDX 10009
378 #define FSCK_BADMDDATAINLN 10010
379 #define FSCK_BADREAD_FBLKMP 10011
380 #define FSCK_BADREAD_FSCKLOG 10012
381 #define FSCK_BADSBAGSIZ 10013
382 #define FSCK_BADSBFJLA 10014
383 #define FSCK_BADSBFJLL 10015
384 #define FSCK_BADSBFSSIZ 10016
385 #define FSCK_BADSBFWSA 10017
386 #define FSCK_BADSBFWSL 10018
387 #define FSCK_BADSBFWSL1 10019
388 #define FSCK_BADSBMGC 10020
389 #define FSCK_BADSBVRSN 10021
390 #define FSCK_BADSBOTHR1 10022
391 #define FSCK_BADSBOTHR2 10023
392 #define FSCK_BADSBOTHR3 10024
393 #define FSCK_BADSBOTHR4 10025
394 #define FSCK_BADSBOTHR5 10026
395 #define FSCK_BADSBOTHR6 10027
396 #define FSCK_BADSBOTHR7 10028
397 #define FSCK_BADSBOTHR8 10029
398 #define FSCK_BADSBOTHR9 10030
399 #define FSCK_BADSBOTHR10 10031
400 #define FSCK_BADSBOTHR11 10032
401 #define FSCK_BADSBOTHR12 10033
402 #define FSCK_BADSBOTHR13 10034
403 #define FSCK_BADWRITE_FSCKLOG 10035
404 #define FSCK_BADWRITE_FBLKMP 10036
405 #define FSCK_BBINOBAD 10037
406 #define FSCK_BLSIZLTLVBSIZ 10038
407 #define FSCK_BMINOBAD 10039
408 #define FSCK_CANT_ALLOC_INOREC 10040
409 #define FSCK_CANT_ALLOC_LSFN 10041
410 #define FSCK_CANT_EXTEND_ROOTDIR 10042
411 #define FSCK_CANTREADEAITEXT1 10043
412 #define FSCK_CANTREADEAITEXT1 10044
413 #define FSCK_CANTREADEAITEXT2 10045
414 #define FSCK_CANTREADEAITEXT3 10046
415 #define FSCK_CANTREADEAITEXT4 10047
416 #define FSCK_CANTREADAGGFSINO 10048
417 #define FSCK_CANTREADBBINO 10049
418 #define FSCK_CANTREADBMINO 10050
419 #define FSCK_CANTREADEA 10051
420 #define FSCK_CANTREADFSEXT 10052
421 #define FSCK_CANTREADFSRTDR 10053
422 #define FSCK_CANTREADLOGINO 10054
423 #define FSCK_CANTREADRECONDNODE 10055
424 #define FSCK_CANTREADRECONDNODE1 10056
425 #define FSCK_CANTREADSELFINO 10057
426 #define FSCK_CANTWRITRECONDNODE 10058
427 #define FSCK_CANTWRITRECONDNODE1 10059
428 #define FSCK_DUPMDBLKREF 10060
429 #define FSCK_FSETEXTBAD 10061
430 #define FSCK_FSRTDRBAD 10062
431 #define FSCK_IAGNOOOAGGBOUNDS 10063
432 #define FSCK_IAGNOOOFSETBOUNDS 10064
433 #define FSCK_INOEXTNOTALLOC 10065
434 #define FSCK_INOINLINECONFLICT1 10066
435 #define FSCK_INOINLINECONFLICT2 10067
436 #define FSCK_INOINLINECONFLICT3 10068
437 #define FSCK_INOINLINECONFLICT4 10069
438 #define FSCK_INOINLINECONFLICT5 10070
439 #define FSCK_LOGINOBAD 10071
440 #define FSCK_RIBADTREE 10072
441 #define FSCK_RIDATAERROR 10073
442 #define FSCK_RINOTDIR 10074
443 #define FSCK_RIUNALLOC 10075
444 #define FSCK_SELFINOBAD 10076
445 #define FSCK_INSUFSTG4RECON 10077
446 #define FSCK_INSUFSTG4RECON1 10078
447 #define FSCK_BLKSNOTAVAILABLE 10079
448 #define FSCK_BADREADTARGET 10080
449 #define FSCK_BADREADTARGET1 10081
450 #define FSCK_BADREADTARGET2 10082

```

```

451 #define FSCK_ENOMEMDBLK1      10083
452 #define FSCK_ENOMEMDBLK2      10084
453 #define FSCK_ENOMEMDBLK3      10085
454 #define FSCK_ENOMEMDBLK4      10086
455 #define FSCK_PARENTNULLIFIED   10087
456 #define FSCK_IOTARGETINJRNLOG  10088
457
458 /* fatal condition return codes */
459
460 #define FSCK_FAILED_SEEK        -10001
461 #define FSCK_FAILED_BADSEEK    -10002
462
463 #define FSCK_FAILED_NODE_BADFLUSH -10003
464 #define FSCK_FAILED_NODE_FLUSH  -10004
465 #define FSCK_FAILED_BADREAD_NODE -10005
466 #define FSCK_FAILED_BADREAD_NODE1 -10006
467 #define FSCK_FAILED_READ_NODE    -10007
468 #define FSCK_FAILED_READ_NODE2   -10008
469 #define FSCK_FAILED_READ_NODE3   -10009
470 #define FSCK_FAILED_READ_NODE4   -10010
471
472 #define FSCK_FAILED_BADREAD_DNODE -10011
473 #define FSCK_FAILED_READ_DNODE   -10012
474
475 #define FSCK_FAILED_INODE_BADFLUSH -10013
476 #define FSCK_FAILED_INODE_FLUSH   -10014
477 #define FSCK_FAILED_BADREAD_INODE -10015
478 #define FSCK_FAILED_BADREAD_INODE1 -10016
479 #define FSCK_FAILED_READ_INODE    -10017
480
481 #define FSCK_FAILED_IAG_BADFLUSH  -10018
482 #define FSCK_FAILED_IAG_FLUSH     -10019
483 #define FSCK_FAILED_BADREAD_IAG   -10020
484 #define FSCK_FAILED_BADREAD1_IAG  -10021
485 #define FSCK_FAILED_READ_IAG      -10022
486 #define FSCK_FAILED_IAG_CORRUPT_PXD -10023
487
488 #define FSCK_FAILED_FBMAP_FLUSH   -10024
489 #define FSCK_FAILED_FBMAP_BADFLUSH -10025
490 #define FSCK_FAILED_BADREAD_FBLKMP -10026
491 #define FSCK_FAILED_READ_FBLKMP   -10027
492 #define FSCK_FAILED_WRITE_FBLKMP  -10028
493 #define FSCK_FAILED_BADWRITE_FBLKMP -10029
494
495 #define FSCK_FAILED_PSBLK_WRITE   -10030
496 #define FSCK_FAILED_SSBLK_WRITE   -10031
497 #define FSCK_FAILED_BTHSBLK_WRITE -10032
498 #define FSCK_FAILED_BTHSBLK_BAD   -10033
499
500 #define FSCK_FAILED_FSSIEXT_READ2 -10034
501 #define FSCK_FAILED_FSRTDIR_READ2 -10035
502 #define FSCK_FAILED_BADBLK_READ2  -10036
503 #define FSCK_FAILED_BMAP_READ2    -10037
504 #define FSCK_FAILED_LOG_READ2     -10038
505 #define FSCK_FAILED_SELF_READ2    -10039
506 #define FSCK_FAILED_SELF_READ3    -10040
507 #define FSCK_FAILED_SELF_READ4    -10041
508 #define FSCK_FAILED_SELF_READ5    -10042
509 #define FSCK_FAILED_SELF_NOWBAD   -10043
510 #define FSCK_FAILED_AGFS_READ2    -10044
511 #define FSCK_FAILED_AGFS_READ3    -10045
512 #define FSCK_FAILED_AGFS_READ4    -10046
513 #define FSCK_FAILED_AGFS_READ5    -10047
514 #define FSCK_FAILED_AGFS_NOWBAD   -10048
515
516 #define FSCK_FAILED_BOTHAITBAD    -10049
517 #define FSCK_FAILED_CANTREADAITEXT1 -10050
518 #define FSCK_FAILED_CANTREADAITEXT2 -10051
519 #define FSCK_FAILED_CANTREADAITEXT3 -10052
520 #define FSCK_FAILED_CANTREADAITEXT4 -10053
521 #define FSCK_FAILED_CANTREADAITEXT5 -10054
522 #define FSCK_FAILED_CANTREADAITEXT6 -10055
523 #define FSCK_FAILED_CANTREADAITEXT7 -10056
524 #define FSCK_FAILED_CANTREADAITEXT8 -10057
525 #define FSCK_FAILED_CANTREADAITEXT9 -10058
526 #define FSCK_FAILED_CANTREADAITEXTA -10059
527 #define FSCK_FAILED_CANTREADAITEXTB -10060
528 #define FSCK_FAILED_CANTREADAITEXTC -10061
529 #define FSCK_FAILED_CANTREADAITEXTD -10062
530 #define FSCK_FAILED_CANTREADAITEXTE -10063
531 #define FSCK_FAILED_CANTREADAITEXTF -10064
532 #define FSCK_FAILED_CANTREADAITEXTG -10065
533 #define FSCK_FAILED_CANTREADAITEXTH -10066
534 #define FSCK_FAILED_CANTREADAITEXTJ -10067
535 #define FSCK_FAILED_CANTREADAITEXTK -10068
536 #define FSCK_FAILED_CANTREADAITCTL -10069
537 #define FSCK_FAILED_CANTREADAITS    -10070
538 #define FSCK_FAILED_CANTREADAIMNOW -10071
539
540 #define FSCK_FAILED_IMPLF_BADFLUSH -10072

```

```
541 #define FSCK_FAILED_IMPLF_FLUSH -10073
542 #define FSCK_FAILED_BADREAD_IMPLF -10074
543 #define FSCK_FAILED_READ_IMPLF -10075
544
545 #define FSCK_FAILED_CANTREAD_DIRNOW -10076
546 #define FSCK_FAILED_DIRGONEBAD -10077
547 #define FSCK_FAILED_DIRGONEBAD2 -10078
548 #define FSCK_FAILED_DIRENTRYGONE -10079
549 #define FSCK_FAILED_DIRENTRYBAD -10080
550
551 #define FSCK_FAILED_MAPCTL_BADFLUSH -10081
552 #define FSCK_FAILED_MAPCTL_FLUSH -10082
553 #define FSCK_FAILED_BADREAD_MAPCTL -10083
554 #define FSCK_FAILED_READ_MAPCTL -10084
555
556 #define FSCK_FAILED_CANTREADBMPCTL -10085
557
558 #define FSCK_FAILED_BMPLV_BADFLUSH -10086
559 #define FSCK_FAILED_BMPLV_FLUSH -10087
560 #define FSCK_FAILED_BADREAD_BMPLV -10088
561 #define FSCK_FAILED_READ_BMPLV -10089
562
563 #define FSCK_FAILED_BMPDM_BADFLUSH -10090
564 #define FSCK_FAILED_BMPDM_FLUSH -10091
565 #define FSCK_FAILED_BADREAD_BMPDM -10092
566 #define FSCK_FAILED_READ_BMPDM -10093
567 #define FSCK_FAILED_DYNSTG_EXHAUST1 -10094
568 #define FSCK_FAILED_DYNSTG_EXHAUST2 -10095
569 #define FSCK_FAILED_DYNSTG_EXHAUST3 -10096
570 #define FSCK_FAILED_DYNSTG_EXHAUST4 -10097
571 #define FSCK_FAILED_DYNSTG_EXHAUST5 -10098
572 #define FSCK_FAILED_DYNSTG_EXHAUST6 -10099
573 #define FSCK_FAILED_DYNSTG_EXHAUST7 -10100
574 #define FSCK_FAILED_DYNSTG_EXHAUST8 -10101
575 #define FSCK_FAILED_DYNSTG_EXHAUST9 -10102
576 #define FSCK_FAILED_DYNSTG_EXHAUSTA -10103
577 #define FSCK_FAILED_REREAD_AGGINO -10104
578
579 /* catastrophic error return codes */
580
581 #define FSCK_INTERNAL_ERROR_1 -11001
582 #define FSCK_INTERNAL_ERROR_2 -11002
583 #define FSCK_INTERNAL_ERROR_3 -11003
584 #define FSCK_INTERNAL_ERROR_4 -11004
585 #define FSCK_INTERNAL_ERROR_5 -11005
586 #define FSCK_INTERNAL_ERROR_6 -11006
587 #define FSCK_INTERNAL_ERROR_7 -11007
588 #define FSCK_INTERNAL_ERROR_8 -11008
589 #define FSCK_INTERNAL_ERROR_9 -11009
590 #define FSCK_INTERNAL_ERROR_10 -11010
591 #define FSCK_INTERNAL_ERROR_11 -11011
592 #define FSCK_INTERNAL_ERROR_12 -11012
593 #define FSCK_INTERNAL_ERROR_13 -11013
594 #define FSCK_INTERNAL_ERROR_14 -11014
595 #define FSCK_INTERNAL_ERROR_15 -11015
596 #define FSCK_INTERNAL_ERROR_16 -11016
597 #define FSCK_INTERNAL_ERROR_17 -11017
598 #define FSCK_INTERNAL_ERROR_18 -11018
599 #define FSCK_INTERNAL_ERROR_19 -11019
600 #define FSCK_INTERNAL_ERROR_20 -11020
601 #define FSCK_INTERNAL_ERROR_21 -11021
602 #define FSCK_INTERNAL_ERROR_22 -11022
603 #define FSCK_INTERNAL_ERROR_23 -11023
604 #define FSCK_INTERNAL_ERROR_24 -11024
605 #define FSCK_INTERNAL_ERROR_25 -11025
606 #define FSCK_INTERNAL_ERROR_26 -11026
607 #define FSCK_INTERNAL_ERROR_27 -11027
608 #define FSCK_INTERNAL_ERROR_28 -11028
609 #define FSCK_INTERNAL_ERROR_29 -11029
610 #define FSCK_INTERNAL_ERROR_30 -11030
611 #define FSCK_INTERNAL_ERROR_31 -11031
612 #define FSCK_INTERNAL_ERROR_32 -11032
613 #define FSCK_INTERNAL_ERROR_33 -11033
614 #define FSCK_INTERNAL_ERROR_34 -11034
615 #define FSCK_INTERNAL_ERROR_35 -11035
616 #define FSCK_INTERNAL_ERROR_36 -11036
617 #define FSCK_INTERNAL_ERROR_37 -11037
618 #define FSCK_INTERNAL_ERROR_38 -11038
619 #define FSCK_INTERNAL_ERROR_39 -11039
620 #define FSCK_INTERNAL_ERROR_40 -11040
621 #define FSCK_INTERNAL_ERROR_41 -11041
622 #define FSCK_INTERNAL_ERROR_42 -11042
623 #define FSCK_INTERNAL_ERROR_43 -11043
624 #define FSCK_INTERNAL_ERROR_44 -11044
625 #define FSCK_INTERNAL_ERROR_45 -11045
626 #define FSCK_INTERNAL_ERROR_46 -11046
627 #define FSCK_INTERNAL_ERROR_47 -11047
628 #define FSCK_INTERNAL_ERROR_48 -11048
629 #define FSCK_INTERNAL_ERROR_49 -11049
630 #define FSCK_INTERNAL_ERROR_50 -11050
```

```
631 #define FCK_INTERNAL_ERROR_51 -11051
632 #define FCK_INTERNAL_ERROR_52 -11052
633 #define FCK_INTERNAL_ERROR_53 -11053
634 #define FCK_INTERNAL_ERROR_54 -11054
635 #define FCK_INTERNAL_ERROR_55 -11055
636 #define FCK_INTERNAL_ERROR_56 -11056
637 #define FCK_INTERNAL_ERROR_57 -11057
638 #define FCK_INTERNAL_ERROR_58 -11058
639 #define FCK_INTERNAL_ERROR_59 -11059
640 #define FCK_INTERNAL_ERROR_60 -11060
641 #define FCK_INTERNAL_ERROR_61 -11061
642 #define FCK_INTERNAL_ERROR_62 -11062
643 #define FCK_INTERNAL_ERROR_63 -11063
644 #define FCK_INTERNAL_ERROR_64 -11064
645 #define FCK_INTERNAL_ERROR_65 -11065
646 #define FCK_INTERNAL_ERROR_66 -11066
647 #define FCK_INTERNAL_ERROR_67 -11067
648 #define FCK_INTERNAL_ERROR_68 -11068
649 #define FCK_INTERNAL_ERROR_69 -11069
650 #define FCK_INTERNAL_ERROR_70 -11070
651 #define FCK_INTERNAL_ERROR_71 -11071
652 #define FCK_INTERNAL_ERROR_72 -11072
653 #define FCK_INTERNAL_ERROR_73 -11073
654 #define FCK_INTERNAL_ERROR_74 -11074
655 #define FCK_INTERNAL_ERROR_75 -11075
656 #define FCK_INTERNAL_ERROR_76 -11076
657 #define FCK_INTERNAL_ERROR_77 -11077
658 #define FCK_INTERNAL_ERROR_78 -11078
659 #define FCK_INTERNAL_ERROR_79 -11079
660 #define FCK_INTERNAL_ERROR_80 -11080
661 #define FCK_INTERNAL_ERROR_81 -11081
662 #define FCK_INTERNAL_ERROR_82 -11082
663 #define FCK_INTERNAL_ERROR_83 -11083
664 #define FCK_INTERNAL_ERROR_84 -11084
665 #define FCK_INTERNAL_ERROR_85 -11085
666 #define FCK_INTERNAL_ERROR_86 -11086
667 #define FCK_INTERNAL_ERROR_87 -11087
668 #define FCK_INTERNAL_ERROR_88 -11088
669 #define FCK_INTERNAL_ERROR_89 -11089
670
671 #endif
```



```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_XFSCCKINT
19 #define H_XFSCCKINT
20
21 #ifndef _JFS_FSCCKDIRE
22 #include "xfscck.h"
23 #endif
24
25 #include "xchkdsk.h"
26 #include "fscckmsgc.h"
27 #include "fscckpfs.h"
28 #include "fscckwsp.h"
29 #include "jfs_endian.h"
30
31 /*
32 ----- functions defined in fscckbmap -----
33 */
34
35 int rebuild_blkall_map(void);
36
37 int verify_blkall_map(void);
38
39 /*
40 ----- functions defined in fscckconn -----
41 */
42
43 int adjust_parents(struct fsck_inode_record *, uint32_t);
44
45 int check_connectedness(void);
46
47 int check_dir_integrity(void);
48
49 int check_link_counts(void);
50
51 /*
52 ----- functions defined in fscckdire -----
53 */
54
55 int fsck_dtDelete(struct dinode *, struct component_name *, uint32_t *);
56
57 int fsck_dtInsert(struct dinode *, struct component_name *, uint32_t *);
58
59 /*
60 ----- functions defined in fscckdtre -----
61 */
62
63 int direntry_add(struct dinode *, uint32_t, UniChar *);
64
65 int direntry_get_inonum(uint32_t, int, UniChar *, int, UniChar *, uint32_t *);
66
67 int direntry_get_objnam(uint32_t, uint32_t, int *, UniChar *);
68
69 int direntry_remove(struct dinode *, uint32_t);
70
71 int dTree_processing(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *, int);
72
73 int dTree_search(struct dinode *, UniChar *, uint32_t, UniChar *,
74                 uint32_t, struct dtslot **, int8_t *, struct fsck_inode_record *);
75
76 int find_first_dir_leaf(struct dinode *, dtpage_t **, int64_t *, int8_t *, int8_t *);
77
78 void init_dir_tree(dtroot_t *);
79
80 int process_valid_dir_data(struct dinode *, uint32_t, struct fsck_inode_record *,
81                           struct fsck_ino_msg_info *, int);
82
83 int reconnect_fs_inodes(void);
84
85 /*
86 ----- functions defined in fscckimap -----
87 */
88
89 int AIS_redundancy_check(void);
90
```

```

91  int AIS_replication(void);
92
93  int rebuild_agg_iamap(void);
94
95  int rebuild_fs_iamaps(void);
96
97  int record_dupchk_inode_extents(void);
98
99  int verify_agg_iamap(void);
100
101  int verify_fs_iamaps(void);
102
103  /*
104  ----- functions defined in fsckino -----
105  */
106
107  #define inode_type_recognized(iptr)\
108      ( ISDIR(iptr->di_mode)  \ \
109        ISREG(iptr->di_mode)  \ \
110        ISLNK(iptr->di_mode)  \ \
111        ISBLK(iptr->di_mode)  \ \
112        ISCHR(iptr->di_mode)  \ \
113        ISFIFO(iptr->di_mode) \ \
114        ISSOCK(iptr->di_mode) )
115
116  int backout_ACL(struct dinode *, struct fsck_inode_record *);
117
118  int backout_EA(struct dinode *, struct fsck_inode_record *);
119
120  int calculate_dasd_used(void);
121
122  int clear_ACL_field(struct fsck_inode_record *, struct dinode *);
123
124  int clear_EA_field(struct fsck_inode_record *, struct dinode *);
125
126  int display_path(uint32_t, int, uint32_t, char *, struct fsck_inode_record *);
127
128  int display_paths(uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
129
130  int first_ref_check_inode(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
131
132  int get_path(uint32_t, uint32_t, char **, struct fsck_inode_record *);
133
134  int in_inode_data_check(struct fsck_inode_record *, struct fsck_ino_msg_info *);
135
136  int inode_is_in_use(struct dinode *, uint32_t);
137
138  int parent_count(struct fsck_inode_record *);
139
140  int record_valid_inode(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
141
142  int release_inode(uint32_t, struct fsck_inode_record *, struct dinode *);
143
144  int unrecord_valid_inode(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
145
146  int validate_ACL(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
147
148  int validate_dasd_used(void);
149
150  int validate_data(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
151
152  int validate_dir_data(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
153
154  int validate_EA(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *);
155
156  int validate_record_fileset_inode(uint32_t, uint32_t, struct dinode *, struct fsck_ino_msg_info *);
157
158  /*
159  ----- functions defined in fsckmeta -----
160  */
161
162  int agg_clean_or_dirty(void);
163
164  int fatal_dup_check(void);
165
166  int first_ref_check_agg_metadata(void);
167
168  int first_ref_check_fixed_metadata(void);
169
170  int first_ref_check_fs_metadata(void);
171
172  int first_ref_check_inode_extents(void);
173
174  int record_fixed_metadata(void);
175
176  int replicate_superblock(void);
177
178  int validate_fs_metadata(void);
179
180  int validate_repair_superblock(void);

```

```
181 int validate_select_agg_inode_table(void);
182
183 /*
184  *----- functions defined in fsckmsgs -----
185 */
186
187 void fsck_send_msg(int, int);
188
189 /*
190  *----- functions defined in fsckpfs -----
191 */
192
193 int ait_node_get(int64_t, xtpage_t *);
194
195 int ait_node_put(int64_t, xtpage_t *);
196
197 int ait_special_read_extl(int);
198
199 int blkmap_find_bit(int64_t, int64_t *, uint32_t *, uint32_t *);
200
201 int blkmap_flush(void);
202
203 int blkmap_get_ctl_page(struct fsck_blk_map_hdr *);
204
205 int blkmap_get_page(int64_t, struct fsck_blk_map_page **);
206
207 int blkmap_put_ctl_page(struct fsck_blk_map_hdr *);
208
209 int blkmap_put_page(int64_t);
210
211 int blktbl_ctl_page_put(struct dbmap *);
212
213 int blktbl_dmap_get(int64_t, struct dmap **);
214
215 int blktbl_dmap_put(struct dmap *);
216
217 int blktbl_dmaps_flush(void);
218
219 int blktbl_Ln_page_get(int8_t, int64_t, struct dmapctl **);
220
221 int blktbl_Ln_page_put(struct dmapctl *);
222
223 int blktbl_Ln_pages_flush(void);
224
225 int close_volume(void);
226
227 int default_volume(void);
228
229 int dnode_get(int64_t, uint32_t, dtpage_t **);
230
231 int ea_get(int64_t, uint32_t, char *, uint32_t *, uint32_t *, int64_t *);
232
233 int fscklog_put_buffer(void);
234
235 int iag_get(int, int, int, int32_t, struct iag **);
236
237 int iag_get_first(int, int, int, struct iag **);
238
239 int iag_get_next(struct iag **);
240
241 int iag_put(struct iag *);
242
243 int iags_flush(void);
244
245 int inodes_flush(void);
246
247 int inode_get(int, int, uint32_t, struct dinode **);
248
249 int inode_get_first_fs(int, uint32_t *, struct dinode **);
250
251 int inode_get_next(uint32_t *, struct dinode **);
252
253 int inode_put(struct dinode *);
254
255 int inotbl_get_ctl_page(int, struct dinomap **);
256
257 int inotbl_put_ctl_page(int, struct dinomap *);
258
259 int mapctl_get(int64_t, void **);
260
261 int mapctl_put(void *);
262
263 int mapctl_flush(void);
264
265 int node_get(int64_t, xtpage_t **);
266
267 int open_volume(char *);
268
269 int readwrite_device(int64_t, uint32_t, uint32_t *, void *, int);
270
```

```
271 int recon_dnode_assign(int64_t, dtpage_t **);
272
273 int recon_dnode_get(int64_t, dtpage_t **);
274
275 int recon_dnode_put(dtpage_t *);
276
277 int recon_dnode_release(dtpage_t *);
278
279 /*
280 ----- functions defined in fsckruns -----
281 */
282
283 void fsck_hbeat_start(void);
284
285 void fsck_hbeat_stop(void);
286
287 /*
288 ----- functions defined in fsckwsp -----
289 */
290
291 int alloc_vlarge_buffer(void);
292
293 int alloc_wrksp(uint32_t, int, int, void **); /* called from both fsck modules
294                                             * and from logredo modules
295                                             */
296
297 int blkall_decrement_owners(int64_t);
298
299 int blkall_increment_owners(int64_t);
300
301 int blkall_ref_check(int64_t, int *);
302
303 int dire_buffer_alloc(dtpage_t **);
304
305 int dire_buffer_release(dtpage_t *);
306
307 int directory_buffers_alloc(void);
308
309 int directory_buffers_release(void);
310
311 int dtreeQ_dequeue(struct dtreeQelem **);
312
313 int dtreeQ_enqueue(struct dtreeQelem *);
314
315 int dtreeQ_get_elem(struct dtreeQelem **);
316
317 int dtreeQ_rel_elem(struct dtreeQelem *);
318
319 int establish_agg_workspace(void);
320
321 int establish_ea_iobuf(void);
322
323 int establish_fs_workspace(void);
324
325 int establish_io_buffers(void);
326
327 int establish_wsp_block_map_ctl(void);
328
329 int extent_record(int64_t, int64_t);
330
331 int extent_unrecord(int64_t, int64_t);
332
333 int fsck_alloc_fsblks(int32_t, int64_t *);
334
335 int fsck_dealloc_fsblks(int32_t, int64_t);
336
337 int fscklog_end(void);
338
339 int fscklog_init(void);
340
341 int fscklog_start(void);
342
343 int get_inode_extension(struct fsck_inode_ext_record **);
344
345 int get_inorecptr(int, int, uint32_t, struct fsck_inode_record **);
346
347 int get_inorecptr_first(int, uint32_t *, struct fsck_inode_record **);
348
349 int get_inorecptr_next(int, uint32_t *, struct fsck_inode_record **);
350
351 int init_agg_record(void);
352
353 int process_extent(struct fsck_inode_record *, uint32_t, int64_t, int8_t,
354                  int8_t, struct fsck_ino_msg_info *, uint32_t *, int8_t *, int);
355
356 int release_inode_extension(struct fsck_inode_ext_record *);
357
358 int release_logredo_allocs(void);
359
360
```

```
361 int temp_inode_buf_alloc(char **);
362
363 int temp_inode_buf_release(char *);
364
365 int temp_node_buf_alloc(char **);
366
367 int temp_node_buf_release(char *);
368
369 int treeQ_dequeue(struct treeQelem **);
370
371 int treeQ_enqueue(struct treeQelem *);
372
373 int treeQ_get_elem(struct treeQelem **);
374
375 int treeQ_rel_elem(struct treeQelem *);
376
377 int treeStack_pop(struct fsck_inode_record **);
378
379 int treeStack_push(struct fsck_inode_record *);
380
381 int workspace_release(void);
382
383 /*
384  ----- functions defined in fsckxtre -----
385 */
386
387 int find_first_leaf(struct dinode *, xtpage_t **, int64_t *, int8_t *, int8_t *);
388
389 int init_xtree_root(struct dinode *);
390
391 int process_valid_data(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *, int);
392
393 int xTree_processing(struct dinode *, uint32_t, struct fsck_inode_record *, struct fsck_ino_msg_info *, int);
394
395 int xTree_search(struct dinode *, int64_t, xad_t **, int8_t *);
396
397 /*
398  ----- functions defined in xchkdsk.c -----
399 */
400
401 void report_readait_error(int, int, int);
402
403 #endif
```

```
1 /*
2  * Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  * This program is free software; you can redistribute it and/or modify
5  * it under the terms of the GNU General Public License as published by
6  * the Free Software Foundation; either version 2 of the License, or
7  * (at your option) any later version.
8  *
9  * This program is distributed in the hope that it will be useful,
10 * but WITHOUT ANY WARRANTY; without even the implied warranty of
11 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 * the GNU General Public License for more details.
13 *
14 * You should have received a copy of the GNU General Public License
15 * along with this program; if not, write to the Free Software
16 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_BTREE
19 #define _H_JFS_BTREE
20 /*
21  * jfs_btree.h: B+-tree
22  *
23  * JFS B+-tree (dtree and xtree) common definitions
24  */
25
26 /*
27  * basic btree page - btpage
28  */
29 struct btpage {
30     int64_t next;          /* 8: right sibling bn */
31     int64_t prev;        /* 8: left sibling bn */
32
33     uint8_t flag;         /* 1: */
34     uint8_t rsrvd[7];    /* 7: type specific */
35     int64_t self;        /* 8: self address */
36
37     uint8_t entry[4064]; /* 4064: */
38 };                       /* (4096) */
39
40 /* btpage flag */
41 #define BT_TYPE      0x07 /* B+-tree index */
42 #define BT_ROOT     0x01 /* root page */
43 #define BT_LEAF     0x02 /* leaf page */
44 #define BT_INTERNAL 0x04 /* internal page */
45 #define BT_RIGHTMOST 0x10 /* rightmost page */
46 #define BT_LEFTMOST 0x20 /* leftmost page */
47 #define BT_SWAPPED  0x80 /* endian swapped when read from disk */
48
49 #endif /* _H_JFS_BTREE */
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_DINODE
19 #define _H_JFS_DINODE
20
21 /*
22 *   jfs_dinode.h: on-disk inode manager
23 *
24 */
25 #include "jfs_types.h"
26 #include "jfs_dtree.h"
27 #include "jfs_xtree.h"
28
29 #define INODESLOTSIZE      128
30 #define L2INODESLOTSIZE   7
31 #define log2INODESIZE     9      /* log2(bytes per dinode) */
32
33 /*
34 *   on-disk inode (struct dinode): 512 bytes
35 *
36 *   * note: align 64-bit fields on 8-byte boundary.
37 */
38 struct dinode {
39     /*
40     *   I. base area (128 bytes)
41     *   -----
42     *
43     *   define generic/POSIX attributes
44     */
45     uint32_t di_inostamp; /* 4: stamp to show inode belongs to fileset */
46     int32_t di_fileset; /* 4: fileset number */
47     uint32_t di_number; /* 4: inode number, aka file serial number */
48     uint32_t di_gen; /* 4: inode generation number */
49
50     pxd_t di_ixpxd; /* 8: inode extent descriptor */
51
52     int64_t di_size; /* 8: size */
53     int64_t di_nblocks; /* 8: number of blocks allocated */
54
55     uint32_t di_nlink; /* 4: number of links to the object */
56
57     uint32_t di_uid; /* 4: user id of owner */
58     uint32_t di_gid; /* 4: group id of owner */
59
60     uint32_t di_mode; /* 4: attribute, format and permission */
61
62     struct timestruc_t di_atime; /* 8: time last data accessed */
63     struct timestruc_t di_ctime; /* 8: time last status changed */
64     struct timestruc_t di_mtime; /* 8: time last data modified */
65     struct timestruc_t di_otime; /* 8: time created */
66
67     dxid_t di_acl; /* 16: acl descriptor */
68
69     dxid_t di_ea; /* 16: ea descriptor */
70
71     int32_t di_next_index; /* 4: Next available dir_table index */
72
73     int32_t di_acltype; /* 4: Type of ACL */
74
75     /*
76     *   Extension Areas.
77     *
78     *   Historically, the inode was partitioned into 4 128-byte areas,
79     *   the last 3 being defined as unions which could have multiple
80     *   uses. The first 96 bytes had been completely unused until
81     *   an index table was added to the directory. It is now more
82     *   useful to describe the last 3/4 of the inode as a single
83     *   union. We would probably be better off redesigning the
84     *   entire structure from scratch, but we don't want to break
85     *   commonality with OS/2's JFS at this time.
86     */
87     union {
88         struct {
89             /*
90             */
91         }
92     }
93 }

```

```

91      * This table contains the information needed to
92      * find a directory entry from a 32-bit index.
93      * If the index is small enough, the table is inline,
94      * otherwise, an x-tree root overlays this table
95      */
96      struct dir_table_slot _table[12];          /* 96: inline */
97
98      dtroot_t             _dtroot;            /* 288: dtree root */
99      } _dir;                                   /* (384) */
100 #define di_dirtable      u._dir._table
101 #define di_dtroot        u._dir._dtroot
102 #define di_parent        di_dtroot.header.idotdot
103 #define di_DASD          di_dtroot.header.DASD
104
105      struct {
106          union {
107              uint8_t _data[96];              /* 96: unused */
108              struct {
109                  void *_imap;               /* 4: unused */
110                  uint32_t _gengen;         /* 4: generator */
111              } _imap;
112          } _ul;                                /* 96: */
113 #define di_gengen        u._file._ul._imap._gengen
114
115          union {
116              uint8_t _xtroot[288];
117              struct {
118                  uint8_t unused[16];        /* 16: */
119                  dxd_t _dxd;               /* 16: */
120                  union {
121                      uint32_t _rdev;        /* 4: */
122                      uint8_t _fastsymlink[128];
123                  } _u;
124                  uint8_t _inlineea[128];
125              } _special;
126          } _u2;
127      } _file;
128 #define di_xtroot        u._file._u2._xtroot
129 #define di_dxd           u._file._u2._special._dxd
130 #define di_btroot        di_xtroot
131 #define di_inlinedata    u._file._u2._special._u
132 #define di_rdev          u._file._u2._special._u._rdev
133 #define di_fastsymlink   u._file._u2._special._u._fastsymlink
134 #define di_inlineea      u._file._u2._special._inlineea
135      } u;
136 };
137
138 /* di_mode */
139 /*
140  * The utilities that are dealing directly with the disk
141  * i-node define the modes as follows. The filesystem itself
142  * should use the standard S_IFMT, etc. defines in stat.h
143  */
144 #define IFMT      0xF000          /* S_IFMT - mask of file type */
145 #define IFDIR     0x4000          /* S_IFDIR - directory */
146 #define IFREG     0x8000          /* S_IFREG - regular file */
147 #define IFLNK     0xA000          /* S_IFLNK - symbolic link */
148 #define IFBLK     0x6000          /* S_IFBLK - block special file */
149 #define IFCHR     0x2000          /* S_IFCHR - character special file */
150 #define IFFIFO    0x1000          /* S_IFFIFO - FIFO */
151 #define IFSOCK    0xC000          /* S_IFSOCK - socket */
152
153 #define ISUID     0x0800          /* S_ISUID - set user id when exec'ing */
154 #define ISGID     0x0400          /* S_ISGID - set group id when exec'ing */
155
156 #define IREAD     0x0100          /* S_IRUSR - read permission */
157 #define IWRITE    0x0080          /* S_IWUSR - write permission */
158 #define IEXEC     0x0040          /* S_IXUSR - execute permission */
159
160 /* extended mode bits (on-disk inode di_mode) */
161 #define IFJOURNAL 0x00010000     /* journalled file */
162 #define ISPARSE   0x00020000     /* sparse file enabled */
163 #define INLINEEA  0x00040000     /* inline EA area free */
164 #define ISWAPFILE 0x00800000     /* file open for pager swap space */
165
166 /* more extended mode bits */
167 #define IDIRECTORY 0x20000000     /* directory (shadow of real bit) */
168
169 #endif /* _H_JFS_DINODE */

```



```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_DMAP
19 #define _H_JFS_DMAP
20
21 #include "jfs_types.h"
22
23 #define BMAPVERSION      1          /* version number */
24 #define TREESIZE         (256+64+16+4+1) /* size of a dmap tree */
25 #define LEAFIND          (64+16+4+1) /* index of 1st leaf of a dmap tree */
26 #define LPERDMAP        256        /* num leaves per dmap tree */
27 #define L2LPERDMAP      8          /* l2 number of leaves per dmap tree */
28 #define DBWORD          32         /* # of blks covered by a map word */
29 #define L2DBWORD        5          /* l2 # of blks covered by a mword */
30 #define BUDMIN          L2DBWORD   /* max free string in a map word */
31 #define BPERDMAP        (LPERDMAP * DBWORD) /* num of blks per dmap */
32 #define L2BPERDMAP      13         /* l2 num of blks per dmap */
33 #define CTLTREESIZE     (1024+256+64+16+4+1) /* size of a dmapctl tree */
34 #define CTLLLEAFIND     (256+64+16+4+1) /* idx of 1st leaf of a dmapctl tree */
35 #define LPERCTL         1024       /* num of leaves per dmapctl tree */
36 #define L2LPERCTL       10         /* l2 num of leaves per dmapctl tree */
37 #define ROOT            0          /* index of the root of a tree */
38 #define NOFREE          ((int8_t) -1) /* no blocks free */
39 #define MAXAG           128        /* max number of allocation groups */
40 #define L2MAXAG         7          /* l2 max num of AG */
41 #define L2MINAGSZ       25         /* l2 of minimum AG size in bytes */
42 #define BMAPBLKNO      0          /* blkno of bmap within the map */
43
44 /*
45 * maximum l2 number of disk blocks at the various dmapctl levels.
46 */
47 #define L2MAXL0SIZE     (L2BPERDMAP + 1 * L2LPERCTL)
48 #define L2MAXL1SIZE     (L2BPERDMAP + 2 * L2LPERCTL)
49 #define L2MAXL2SIZE     (L2BPERDMAP + 3 * L2LPERCTL)
50
51 /*
52 * maximum number of disk blocks at the various dmapctl levels.
53 */
54 #define MAXL0SIZE       ((int64_t)1 << L2MAXL0SIZE)
55 #define MAXL1SIZE       ((int64_t)1 << L2MAXL1SIZE)
56 #define MAXL2SIZE       ((int64_t)1 << L2MAXL2SIZE)
57
58 #define MAXMAPSIZE      MAXL2SIZE   /* maximum aggregate map size */
59
60 /*
61 * determine the maximum free string for four (lower level) nodes
62 * of the tree.
63 */
64 #define TREEMAX(cp)     \
65     ((signed char)(MAX(MAX(*(cp),*((cp)+1)), \
66     MAX(*(cp)+2),*((cp)+3))))
67
68 /*
69 * convert disk block number to the logical block number of the dmap
70 * describing the disk block. s is the log2(number of logical blocks per page)
71 *
72 * The calculation figures out how many logical pages are in front of the dmap.
73 * - the number of dmaps preceding it
74 * - the number of L0 pages preceding its L0 page
75 * - the number of L1 pages preceding its L1 page
76 * - 3 is added to account for the L2, L1, and L0 page for this dmap
77 * - 1 is added to account for the control page of the map.
78 */
79 #define BLKTODMAP(b,s)  \
80     (((b) >> 13) + ((b) >> 23) + ((b) >> 33) + 3 + 1) << (s)
81
82 /*
83 * convert disk block number to the logical block number of the LEVEL 0
84 * dmapctl describing the disk block. s is the log2(number of logical blocks
85 * per page)
86 *
87 * The calculation figures out how many logical pages are in front of the L0.
88 * - the number of dmap pages preceding it
89 * - the number of L0 pages preceding it
90 * - the number of L1 pages preceding its L1 page

```

```

91 *      - 2 is added to account for the L2, and L1 page for this L0
92 *      - 1 is added to account for the control page of the map.
93 */
94 #define BLKTOL0(b,s) \
95     (((b) >> 23) << 10) + ((b) >> 23) + ((b) >> 33) + 2 + 1 << (s)
96
97 /*
98 * convert disk block number to the logical block number of the LEVEL 1
99 * dmapctl describing the disk block. s is the log2(number of logical blocks
100 * per page)
101 *
102 * The calculation figures out how many logical pages are in front of the L1.
103 * - the number of dmap pages preceding it
104 * - the number of L0 pages preceding it
105 * - the number of L1 pages preceding it
106 * - 1 is added to account for the L2 page
107 * - 1 is added to account for the control page of the map.
108 */
109 #define BLKTOL1(b,s) \
110     (((b) >> 33) << 20) + ((b) >> 33) << 10) + ((b) >> 33) + 1 + 1 << (s)
111
112 /*
113 * convert disk block number to the logical block number of the dmapctl
114 * at the specified level which describes the disk block.
115 */
116 #define BLKTOCTL(b,s,l) \
117     (((l) == 2) ? 1 : ((l) == 1) ? BLKTOL1(b),(s)) : BLKTOL0(b),(s))
118
119 /*
120 * convert aggregate map size to the zero origin dmapctl level of the
121 * top dmapctl.
122 */
123 #define BMAPSZTOLEV(size) \
124     ((size) <= MAXL0SIZE) ? 0 : ((size) <= MAXL1SIZE) ? 1 : 2)
125
126 /* convert disk block number to allocation group number.
127 */
128 #define BLKTOAG(b,sb) ((b) >> ((sb)->s_jfs_bmap->db_agl2size))
129
130 /* convert allocation group number to starting disk block
131 * number.
132 */
133 #define AGTOBLK(a,ip) \
134     ((int64_t)(a) << ((ip)->i_sb->s_jfs_bmap->db_agl2size))
135
136 /*
137 *      dmap summary tree
138 *
139 * struct dmaptree must be consistent with struct dmapctl.
140 */
141 struct dmaptree {
142     int32_t  nleafs;           /* 4: number of tree leafs      */
143     int32_t  l2nleafs;       /* 4: l2 number of tree leafs   */
144     int32_t  leafidx;        /* 4: index of first tree leaf  */
145     int32_t  height;         /* 4: height of the tree        */
146     int8_t   budmin;         /* 1: min l2 tree leaf value to combine */
147     int8_t   stree[TREESIZE]; /* TREESIZE: tree               */
148     uint8_t  pad[2];         /* 2: pad to word boundary      */
149 }; /* - 360 - */
150
151 /*
152 *      dmap page per 8K blocks bitmap
153 */
154 struct dmap {
155     int32_t  nblocks;        /* 4: num blks covered by this dmap */
156     int32_t  nfree;          /* 4: num of free blks in this dmap */
157     int64_t  start;          /* 8: starting blkno for this dmap */
158     struct dmaptree tree;    /* 360: dmap tree                */
159     uint8_t  pad[1672];      /* 1672: pad to 2048 bytes        */
160     uint32_t wmap[LPERDMAP]; /* 1024: bits of the working map   */
161     uint32_t pmap[LPERDMAP]; /* 1024: bits of the persistent map */
162 }; /* - 4096 - */
163
164 /*
165 *      disk map control page per level.
166 *
167 * struct dmapctl must be consistent with struct dmaptree.
168 */
169 struct dmapctl {
170     int32_t  nleafs;           /* 4: number of tree leafs      */
171     int32_t  l2nleafs;       /* 4: l2 number of tree leafs   */
172     int32_t  leafidx;        /* 4: index of the first tree leaf */
173     int32_t  height;         /* 4: height of tree            */
174     int8_t   budmin;         /* 1: minimum l2 tree leaf value */
175     int8_t   stree[CTLTREESIZE]; /* CTLTREESIZE: dmapctl tree */
176     uint8_t  pad[2714];      /* 2714: pad to 4096            */
177 }; /* - 4096 - */
178
179 /*
180 *      common definition for dmaptree within dmap and dmapctl

```

```

181  */
182  typedef union {
183      struct dmaptree t1;
184      struct dmapctl t2;
185  } dmtree_t;
186
187  /* macros for accessing fields within dmtree_t */
188  #define dmt_nleafs      t1.nleafs
189  #define dmt_l2nleafs   t1.l2nleafs
190  #define dmt_leafidx    t1.leafidx
191  #define dmt_height     t1.height
192  #define dmt_budmin     t1.budmin
193  #define dmt_stree      t1.stree
194
195  /*
196   *      on-disk aggregate disk allocation map descriptor.
197   */
198  struct dbmap {
199      int64_t  dn_mapsize;           /* 8: number of blocks in aggregate */
200      int64_t  dn_nfree;           /* 8: num free blks in aggregate map */
201      int32_t  dn_l2nbperpage;     /* 4: number of blks per page */
202      int32_t  dn_numag;          /* 4: total number of ags */
203      int32_t  dn_maxlevel;       /* 4: number of active ags */
204      int32_t  dn_maxag;         /* 4: max active alloc group number */
205      int32_t  dn_agpref;        /* 4: preferred alloc group (hint) */
206      int32_t  dn_aglevel;       /* 4: dmapctl level holding the AG */
207      int32_t  dn_agheight;      /* 4: height in dmapctl of the AG */
208      int32_t  dn_agwidth;       /* 4: width in dmapctl of the AG */
209      int32_t  dn_agstart;       /* 4: start tree index at AG height */
210      int32_t  dn_agl2size;      /* 4: l2 num of blks per alloc group */
211      int64_t  dn_agfree[MAXAG]; /* 8*MAXAG: per AG free count */
212      int64_t  dn_agsize;        /* 8: num of blks per alloc group */
213      int8_t   dn_maxfreebud;    /* 1: max free buddy system */
214      uint8_t  pad[3007];       /* 3007: pad to 4096 */
215  };
216
217  #endif

```

```

/* _H_JFS_DMAP */

```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_DTREE
19 #define _H_JFS_DTREE
20
21 /*
22 *   jfs_dtree.h: directory B+-tree manager
23 */
24 #include "jfs_btrec.h"
25
26 typedef union {
27     struct {
28         int          tid;
29         struct inode *ip;
30         uint32_t     ino;
31     } leaf;
32     pxd_t xd;
33 } ddata_t;
34
35 /*
36 *   entry segment/slot
37 *
38 *   * an entry consists of type dependent head/only segment/slot and
39 *   * additional segments/slots linked vi next field;
40 *   * N.B. last/only segment of entry is terminated by next = -1;
41 */
42 /*
43 *   directory page slot
44 */
45 struct dtslot {
46     int8_t  next;          /* 1: */
47     int8_t  cnt;          /* 1: */
48     UniChar name[15];     /* 30: */
49 };                        /* (32) */
50
51
52 #define DATASLOTSIZE 16
53 #define L2DATASLOTSIZE 4
54 #define DTSLOTSIZE 32
55 #define L2DTSLOTSIZE 5
56 #define DTSLOTHDRSIZE 2
57 #define DTSLOTDATASIZE 30
58 #define DTSLOTDATALEN 15
59
60
61 /*
62 *   internal node entry head/only segment
63 */
64 struct identry {
65     pxd_t  xd;            /* 8: child extent descriptor */
66
67     int8_t next;          /* 1: */
68     uint8_t namlen;       /* 1: */
69     UniChar name[11];     /* 22: 2-byte aligned */
70 };                        /* (32) */
71
72 #define DTIHDRSIZE 10
73 #define DTIHDRDATALEN 11
74
75 /* compute number of slots for entry */
76 #define NDTINTERNAL(klen) ( ((4 + (klen)) + (15 - 1)) / 15 )
77
78 /*
79 *   leaf node entry head/only segment
80 *
81 *   * For legacy filesystems, name contains 13 unichars -- no index field
82 */
83 struct ldentry {
84     uint32_t inumber;     /* 4: 4-byte aligned */
85     int8_t  next;        /* 1: */
86     uint8_t namlen;       /* 1: */
87     UniChar name[11];     /* 22: 2-byte aligned */
88     uint32_t index;       /* 4: index into dir_table */
89 };                        /* (32) */
90

```

```

91
92 #define DTLHDRSIZE 6
93 #define DTLHDRDATALEN_LEGACY 13 /* Old (OS/2) format */
94 #define DTLHDRDATALEN 11
95
96 /*
97  * dir_table used for directory traversal during readdir
98  */
99
100 /*
101  * Keep persistent index for directory entries
102  */
103 #define DO_INDEX(INODE) ((INODE)->i_sb->s_jfs_mntflag & JFS_DIR_INDEX)
104
105 /*
106  * Maximum entry in inline directory table
107  */
108 #define MAX_INLINE_DIRTABLE_ENTRY 13
109
110 struct dir_table_slot {
111     uint8_t rsvrd; /* 1: */
112     uint8_t flag; /* 1: 0 if free */
113     uint8_t slot; /* 1: slot within leaf page of entry */
114     uint8_t addr1; /* 1: upper 8 bits of leaf page address */
115     uint32_t addr2; /* 4: lower 32 bits of leaf page address -OR- index
116                    of next entry when this entry was deleted */
117 }; /* (8) */
118
119 /*
120  * flag values
121  */
122 #define DIR_INDEX_VALID 1
123 #define DIR_INDEX_FREE 0
124
125 /* compute number of slots for entry */
126 #define NDTLEAF_LEGACY(klen) ((2 + (klen)) + (15 - 1)) / 15 )
127 #define NDTLEAF NDTINTERNAL
128
129 /*
130  * directory root page (in-line in on-disk inode):
131  *
132  * cf. dtpage_t below.
133  */
134 #define
135 typedef union {
136     struct {
137         struct dasd DASD; /* 16: DASD limit/usage info F226941 */
138
139         uint8_t flag; /* 1: */
140         int8_t nextindex; /* 1: next free entry in stbl */
141         int8_t freecnt; /* 1: free count */
142         int8_t freelist; /* 1: freelist header */
143
144         uint32_t idotdot; /* 4: parent inode number */
145
146         int8_t stbl[8]; /* 8: sorted entry index table */
147     } header; /* (32) */
148
149     struct dtslot slot[9];
150 } dtroot_t;
151
152 #define DTROOTMAXSLOT 9
153
154 /*
155  * directory regular page:
156  *
157  * entry slot array of 32 byte slot
158  *
159  * sorted entry slot index table (stbl):
160  * contiguous slots at slot specified by stblindex,
161  * 1-byte per entry
162  * 512 byte block: 16 entry tbl (1 slot)
163  * 1024 byte block: 32 entry tbl (1 slot)
164  * 2048 byte block: 64 entry tbl (2 slot)
165  * 4096 byte block: 128 entry tbl (4 slot)
166  *
167  * data area:
168  * 512 byte block: 16 - 2 = 14 slot
169  * 1024 byte block: 32 - 2 = 30 slot
170  * 2048 byte block: 64 - 3 = 61 slot
171  * 4096 byte block: 128 - 5 = 123 slot
172  *
173  * N.B. index is 0-based; index fields refer to slot index
174  * except nextindex which refers to entry index in stbl;
175  * end of entry slot list or freelist is marked with -1.
176  */
177 typedef union {
178     struct {
179         int64_t next; /* 8: next sibling */
180         int64_t prev; /* 8: previous sibling */

```

```

181         uint8_t  flag;          /* 1: */
182         int8_t   nextindex;     /* 1: next entry index in stbl */
183         int8_t   freecnt;       /* 1: */
184         int8_t   freelist;      /* 1: slot index of head of freelist */
185
186
187         uint8_t  maxslot;       /* 1: number of slots in page slot[] */
188         int8_t   stblindex;     /* 1: slot index of start of stbl */
189         uint8_t  rsvrd[2];      /* 2: */
190
191         pxd_t    self;          /* 8: self pxd */
192     } header;                  /* (32) */
193
194     struct dtslot slot[128];
195 } dtpage_t;
196
197 #define DTPAGEMAXSLOT        128
198
199 #define DT8THPGNODEBYTES     512
200 #define DT8THPGNODETSLOTS   1
201 #define DT8THPGNODESLOTS    16
202
203 #define DTQTRPGNODEBYTES     1024
204 #define DTQTRPGNODETSLOTS   1
205 #define DTQTRPGNODESLOTS    32
206
207 #define DTHALFPGNODEBYTES    2048
208 #define DTHALFPGNODETSLOTS  2
209 #define DTHALFPGNODESLOTS   64
210
211 #define DTFULLPGNODEBYTES    4096
212 #define DTFULLPGNODETSLOTS  4
213 #define DTFULLPGNODESLOTS   128
214
215 #define DTENTRYSTART        1
216
217 /* get sorted entry table of the page */
218 #define DT_GETSTBL(p) ( ((p)->header.flag & BT_ROOT) ? \
219     ((dtroot_t *) (p))->header.stbl : \
220     (int8_t *) &(p)->slot[(p)->header.stblindex] )
221
222 /*
223  * Flags for dtSearch
224  */
225 #define JFS_CREATE 1
226 #define JFS_LOOKUP 2
227 #define JFS_REMOVE 3
228 #define JFS_RENAME 4
229
230 #endif                          /* !_H_JFS_DTREE */

```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_FILSYS
19 #define _H_JFS_FILSYS
20
21 /*
22 *   jfs_filsys.h
23 *
24 *   file system (implementation-dependent) constants
25 *
26 *   refer to <limits.h> for system wide implementation-dependent constants
27 */
28
29 #include "jfs_types.h"
30
31 /*
32 *   file system option (superblock flag)
33 */
34 /* platform option (conditional compilation) */
35 #define JFS_AIX          0x80000000 /* AIX support */
36 /* POSIX name/directory support */
37
38 #define JFS_OS2          0x40000000 /* OS/2 support */
39 /* case-insensitive name/directory support */
40
41 #define JFS_DFS          0x20000000 /* DCE DFS LFS support */
42
43 #define JFS_LINUX        0x10000000 /* Linux support */
44 /* case-sensitive name/directory support */
45
46 /* directory option */
47 #define JFS_UNICODE      0x00000001 /* unicode name */
48
49 /* commit option */
50 #define JFS_COMMIT       0x00000f00 /* commit option mask */
51 #define JFS_GROUPCOMMIT 0x00000100 /* group (of 1) commit */
52 #define JFS_LAZYCOMMIT  0x00000200 /* lazy commit */
53 #define JFS_TMPFS        0x00000400 /* temporary file system -
54 * do not log/commit:
55 */
56
57 /* log logical volume option */
58 #define JFS_INLINELOG    0x00000800 /* inline log within file system */
59 #define JFS_INLINEMOVE  0x00001000 /* inline log being moved */
60
61 /* Secondary aggregate inode table */
62 #define JFS_BAD_SAIT     0x00010000 /* current secondary ait is bad */
63
64 /* sparse regular file support */
65 #define JFS_SPARSE       0x00020000 /* sparse regular file */
66
67 /* DASD Limits F226941 */
68 #define JFS_DASD_ENABLED 0x00040000 /* DASD limits enabled */
69 #define JFS_DASD_PRIME   0x00080000 /* Prime DASD usage on boot */
70
71 /* Directory index */
72 #define JFS_DIR_INDEX    0x00200000 /* Persistent index for
73 /* directory entries */
74
75
76 /*
77 *   buffer cache configuration
78 */
79 /* page size */
80 #ifdef PSIZE
81 #undef PSIZE
82 #endif
83 #define PSIZE            4096 /* page size (in byte) */
84 #define L2PSIZE          12 /* log2(PSIZE) */
85 #define POFFSET          4095 /* offset within page */
86
87 /* buffer page size */
88 #define BPSIZE           PSIZE
89
90 /*

```

```

91  *      fs fundamental size
92  *
93  * PSIZE >= file system block size >= PBSIZE >= DISIZE
94  */
95  #define PBSIZE      512      /* physical block size (in byte) */
96  #define L2PBSIZE   9        /* log2(PBSIZE) */
97
98  #define DISIZE     512      /* on-disk inode size (in byte) */
99  #define L2DISIZE   9        /* log2(DISIZE) */
100
101  #define IDATASIZE  256      /* inode inline data size */
102  #define IXATTRSIZE 128      /* inode inline extended attribute size */
103
104  #define XTPAGE_SIZE 4096
105  #define log2_PAGESIZE 12
106
107  #define IAG_SIZE    4096
108  #define IAG_EXTENT_SIZE 4096
109  #define INOSPERIAG  4096      /* number of disk inodes per iag */
110  #define L2INOSPERIAG 12      /* 12 number of disk inodes per iag */
111  #define INOSPEREXT  32      /* number of disk inode per extent */
112  #define L2INOSPEREXT 5        /* 12 number of disk inode per extent */
113  #define IXSIZE      (DISIZE * INOSPEREXT) /* inode extent size */
114  #define INOSPERPAGE 8        /* number of disk inodes per 4K page */
115  #define L2INOSPERPAGE 3      /* log2(INOSPERPAGE) */
116
117  #define IAGFREELIST_LWM 64
118
119  #define INODE_EXTENT_SIZE IXSIZE /* inode extent size */
120  #define NUM_INODE_PER_EXTENT INOSPEREXT
121  #define NUM_INODE_PER_IAG INOSPERIAG
122
123  #define MINBLOCKSIZE 512
124  #define MAXBLOCKSIZE 4096
125  #define MAXFILESIZE ((int64_t)1 << 52)
126
127  #define JFS_LINK_MAX 65535 /* nlink_t is unsigned short */
128
129  /* Minimum number of bytes supported for a JFS partition */
130  #define MINJFS (0x1000000)
131  #define MINJFSTEXT "16"
132
133  /*
134  * fixed physical block address (physical block size = 512 byte)
135  *
136  * NOTE: since we can't guarantee a physical block size of 512 bytes the use of
137  * these macros should be removed and the byte offset macros used instead.
138  */
139  #define SUPER1_B 64 /* primary superblock */
140  #define AIMAP_B (SUPER1_B + 8) /* 1st extent of aggregate inode map */
141  #define AITBL_B (AIMAP_B + 16) /*
142  * 1st extent of aggregate inode table
143  */
144  #define SUPER2_B (AITBL_B + 32) /* 2ndary superblock pbn */
145  #define BMAP_B (SUPER2_B + 8) /* block allocation map */
146
147  /*
148  * SIZE_OF_SUPER defines the total amount of space reserved on disk for the
149  * superblock. This is not the same as the superblock structure, since all of
150  * this space is not currently being used.
151  */
152  #define SIZE_OF_SUPER PSIZE
153
154  /*
155  * SIZE_OF_AG_TABLE defines the amount of space reserved to hold the AG table
156  */
157  #define SIZE_OF_AG_TABLE PSIZE
158
159  /*
160  * SIZE_OF_MAP_PAGE defines the amount of disk space reserved for each page of
161  * the inode allocation map (to hold iag)
162  */
163  #define SIZE_OF_MAP_PAGE PSIZE
164
165  /*
166  * fixed byte offset address
167  */
168  #define SUPER1_OFF 0x8000 /* primary superblock */
169  #define AIMAP_OFF (SUPER1_OFF + SIZE_OF_SUPER)
170  /*
171  * Control page of aggregate inode map
172  * followed by 1st extent of map
173  */
174  #define AITBL_OFF (AIMAP_OFF + (SIZE_OF_MAP_PAGE << 1))
175  /*
176  * 1st extent of aggregate inode table
177  */
178  #define SUPER2_OFF (AITBL_OFF + INODE_EXTENT_SIZE)
179  /*
180  * secondary superblock

```



```

181                                     */
182 #define BMAP_OFF          (SUPER2_OFF + SIZE_OF_SUPER)
183                                     /*
184                                     * block allocation map
185                                     */
186
187 /*
188 * The following macro is used to indicate the number of reserved disk blocks at
189 * the front of an aggregate, in terms of physical blocks. This value is
190 * currently defined to be 32K. This turns out to be the same as the primary
191 * superblock's address, since it directly follows the reserved blocks.
192 */
193 #define AGGR_RSVD_BLOCKS      SUPER1_B
194
195 /*
196 * The following macro is used to indicate the number of reserved bytes at the
197 * front of an aggregate. This value is currently defined to be 32K. This
198 * turns out to be the same as the primary superblock's byte offset, since it
199 * directly follows the reserved blocks.
200 */
201 #define AGGR_RSVD_BYTES      SUPER1_OFF
202
203 /*
204 * The following macro defines the byte offset for the first inode extent in
205 * the aggregate inode table. This allows us to find the self inode to find the
206 * rest of the table. Currently this value is 44K.
207 */
208 #define AGGR_INODE_TABLE_START  AITBL_OFF
209
210 /*
211 *      fixed reserved inode number
212 */
213 /* aggregate inode */
214 #define AGGR_RESERVED_I      0      /* aggregate inode (reserved) */
215 #define AGGREGATE_I          1      /* aggregate inode map inode */
216 #define BMAP_I                2      /* aggregate block allocation map inode */
217 #define LOG_I                 3      /* aggregate inline log inode */
218 #define BADBLOCK_I           4      /* aggregate bad block inode */
219 #define FILESYSTEM_I         16     /* 1st/only fileset inode in ait:
220                                     * fileset inode map inode
221                                     */
222
223 /* per fileset inode */
224 #define FILESET_RSVD_I        0      /* fileset inode (reserved) */
225 #define FILESET_EXT_I         1      /* fileset inode extension */
226 #define ROOT_I                2      /* fileset root inode */
227 #define ACL_I                 3      /* fileset ACL inode */
228
229 #define FILESET_OBJECT_I      4      /* the first fileset inode available for a file
230                                     * or directory or link...
231                                     */
232 #define FIRST_FILESET_INO     16     /* the first aggregate inode which describes
233                                     * an inode. (To fsck this is also the first
234                                     * inode in part 2 of the agg inode table.)
235                                     */
236
237 /*
238 *      directory configuration
239 */
240 #define JFS_NAME_MAX          255
241 #define JFS_PATH_MAX          BPSIZE
242
243
244 /*
245 *      file system state (superblock state)
246 */
247 #define FM_CLEAN              0x00000000 /* file system is unmounted and clean */
248 #define FM_MOUNT              0x00000001 /* file system is mounted cleanly */
249 #define FM_DIRTY              0x00000002 /* file system was not unmounted and clean
250                                     * when mounted or
251                                     * commit failure occurred while being mounted:
252                                     * fsck() must be run to repair
253                                     */
254 #define FM_LOGREDO            0x00000004 /* log based recovery (logredo()) failed:
255                                     * fsck() must be run to repair
256                                     */
257 #define FM_EXTENDFS           0x00000008 /* file system extendfs() in progress */
258
259 #endif                          /* _H_JFS_FILSYS */

```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_IMAP
19 #define _H_JFS_IMAP
20
21 /*
22 *   jfs_imap.h: disk inode manager
23 */
24
25 #define EXTSPERIAG      128 /* number of disk inode extent per iag */
26 #define IMAPBLKNO      0 /* lblkno of dinomap within inode map */
27 #define SMAPSZ         4 /* number of words per summary map */
28 #define EXTSPERSUM     32 /* number of extents per summary map entry */
29 #define L2EXTSPERSUM  5 /* l2 number of extents per summary map */
30 #define PGSPERIEXT     4 /* number of 4K pages per dinode extent */
31 #define MAXIAGS        ((1<<20)-1) /* maximum number of iags */
32 #define MAXAG          128 /* maximum number of allocation groups */
33
34 #define AMAPSIZE       512 /* bytes in the IAG allocation maps */
35 #define SMAPSIZE       16 /* bytes in the IAG summary maps */
36
37 /* convert inode number to iag number */
38 #define INOTOIAG(ino)  ((ino) >> L2INOSPERIAG)
39
40 /* convert iag number to logical block number of the iag page */
41 #define IAGTOLBLK(iagno,l2nbperpg)  (((iagno) + 1) << (l2nbperpg))
42
43 /* get the starting block number of the 4K page of an inode extent
44 * that contains ino.
45 */
46 #define INOPBLK(pxd,ino,l2nbperpg)  (addressPXD((pxd)) + \
47 (((ino) & (INOSPEREXT-1)) >> L2INOSPERPAGE) << (l2nbperpg))
48
49 /*
50 *   inode allocation map:
51 *
52 *   inode allocation map consists of
53 *   . the inode map control page and
54 *   . inode allocation group pages (per 4096 inodes)
55 *   which are addressed by standard JFS xtree.
56 */
57 /*
58 *   inode allocation group page (per 4096 inodes of an AG)
59 */
60 struct iag {
61     int64_t  agstart; /* 8: starting block of ag */
62     int32_t  iagnum; /* 4: inode allocation group number */
63     int32_t  inofreefwd; /* 4: ag inode free list forward */
64     int32_t  inofreeback; /* 4: ag inode free list back */
65     int32_t  extfreefwd; /* 4: ag inode extent free list forward */
66     int32_t  extfreeback; /* 4: ag inode extent free list back */
67     int32_t  iagfree; /* 4: iag free list */
68
69     /* summary map: 1 bit per inode extent */
70     int32_t  inosmap[SMAPSZ]; /* 16: sum map of mapwords w/ free inodes;
71 * note: this indicates free and backed
72 * inodes, if the extent is not backed the
73 * value will be 1. if the extent is
74 * backed but all inodes are being used the
75 * value will be 1. if the extent is
76 * backed but at least one of the inodes is
77 * free the value will be 0.
78 */
79     int32_t  extsmap[SMAPSZ]; /* 16: sum map of mapwords w/ free extents */
80     int32_t  nfreeinos; /* 4: number of free inodes */
81     int32_t  nfreeexts; /* 4: number of free extents */
82     /* (72) */
83     uint8_t  pad[1976]; /* 1976: pad to 2048 bytes */
84     /* allocation bit map: 1 bit per inode (0 - free, 1 - allocated) */
85     uint32_t  wmap[EXTSPERIAG]; /* 512: working allocation map */
86     uint32_t  pmap[EXTSPERIAG]; /* 512: persistent allocation map */
87     pxd_t  inoext[EXTSPERIAG]; /* 1024: inode extent addresses */
88 }; /* (4096) */
89
90 /*

```

```
91  *      per AG control information (in inode map control page)
92  */
93  struct iagctl {
94      int32_t  inofree;          /* 4: free inode list anchor      */
95      int32_t  extfree;         /* 4: free extent list anchor     */
96      int32_t  numinos;         /* 4: number of backed inodes     */
97      int32_t  numfree;         /* 4: number of free inodes       */
98  };                               /* (16) */
99
100 /*
101  *      per fileset/aggregate inode map control page
102  */
103  struct dinomap {
104      int32_t  in_freeiag;       /* 4: free iag list anchor        */
105      int32_t  in_nextiag;       /* 4: next free iag number        */
106      int32_t  in_numinos;       /* 4: num of backed inodes        */
107      int32_t  in_numfree;       /* 4: num of free backed inodes   */
108      int32_t  in_nbperext;      /* 4: num of blocks per inode extent */
109      int32_t  in_l2nbperext;    /* 4: l2 of in_nbperext          */
110      int32_t  in_diskblock;     /* 4: for standalone test driver  */
111      int32_t  in_maxag;         /* 4: for standalone test driver  */
112      uint8_t  pad[2016];        /* 2016: pad to 2048             */
113      struct iagctl in_agctl[MAXAG]; /* 2048: AG control information */
114  };                               /* (4096) */
115
116 #endif                               /* _H_JFS_IMAP */
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_SUPERBLOCK
19 #define _H_JFS_SUPERBLOCK
20
21 #ifndef _H_JFS_LOGMGR
22 #include <uuid/uuid.h>
23 #endif
24
25 /*
26 *       jfs_superblock.h
27 */
28
29 /*
30 * make the magic number something a human could read
31 */
32 #define JFS_MAGIC      "JFS1"   /* Magic word: JFS 1 */
33
34 #define JFS_VERSION    2        /* Version number: Version 2 */
35
36 #define LV_NAME_SIZE   11       /* MUST BE 11 for OS/2 boot sector */
37
38 /*
39 *       aggregate superblock
40 */
41 struct superblock {
42     char s_magic[4];           /* 4: magic number */
43     uint32_t s_version;       /* 4: version number */
44
45     int64_t s_size;           /* 8: aggregate size in hardware/LVM blocks;
46                               * VFS: number of blocks
47                               */
48     int32_t s_bsize;         /* 4: aggregate block size in bytes;
49                               * VFS: fragment size
50                               */
51     int16_t s_l2bsize;       /* 2: log2 of s_bsize */
52     int16_t s_l2bfactor;     /* 2: log2(s_bsize/hardware block size) */
53     int32_t s_pbsize;        /* 4: hardware/LVM block size in bytes */
54     int16_t s_l2pbsize;      /* 2: log2 of s_pbsize */
55     int16_t pad;             /* 2: padding necessary for alignment */
56
57     uint32_t s_agsize;       /* 4: allocation group size in aggr. blocks */
58
59     uint32_t s_flag;         /* 4: aggregate attributes:
60                               * see jfs_filsys.h
61                               */
62     uint32_t s_state;        /* 4: mount/unmount/recovery state:
63                               * see jfs_filsys.h
64                               */
65     int32_t s_compress;      /* 4: > 0 if data compression */
66
67     pxd_t s_ait2;           /* 8: first extent of secondary
68                               * aggregate inode table
69                               */
70
71     pxd_t s_aim2;           /* 8: first extent of secondary
72                               * aggregate inode map
73                               */
74     uint32_t s_logdev;       /* 4: device address of log */
75     int32_t s_logserial;     /* 4: log serial number at aggregate mount */
76     pxd_t s_logpxd;         /* 8: inline log extent */
77
78     pxd_t s_fsckpxd;        /* 8: inline fsck work space extent */
79
80     struct timestruc_t s_time; /* 8: time last updated */
81
82     int32_t s_fsckloglen;    /* 4: Number of filesystem blocks reserved for
83                               * the fsck service log.
84                               * N.B. These blocks are divided among the
85                               * versions kept. This is not a per
86                               * version size.
87                               * N.B. These blocks are included in the
88                               * length field of s_fsckpxd.
89                               */
90     int8_t s_fscklog;       /* 1: which fsck service log is most recent

```

```
91          *      0 => no service log data yet
92          *      1 => the first one
93          *      2 => the 2nd one
94          */
95      char s_fpack[11]; /* 11: file system volume name
96          *      N.B. This must be 11 bytes to
97          *      conform with the OS/2 BootSector
98          *      requirements
99          *      Only used when s_version is 1
100         */
101
102     /* extendfs() parameter under s_state & FM_EXTENDFS */
103     int64_t s_xsize; /* 8: extendfs s_size */
104     pxd_t s_xfsckpxd; /* 8: extendfs fsckpxd */
105     pxd_t s_xlogpxd; /* 8: extendfs logpxd */
106     /* - 128 byte boundary - */
107
108     uuid_t s_uuid; /* 16: 128-bit uuid for volume */
109     char s_label[16]; /* 16: volume label */
110     uuid_t s_loguuid; /* 16: 128-bit uuid for log device */
111 };
112
113 #endif /* _H_JFS_SUPERBLOCK */
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_XTREE
19 #define _H_JFS_XTREE
20
21 /*
22 *   jfs_xtree.h: extent allocation descriptor B+-tree manager
23 */
24
25 #include "jfs_btree.h"
26
27 /*
28 *   extent allocation descriptor (xad)
29 */
30 typedef struct xad {
31     unsigned flag:8;          /* 1: flag */
32     unsigned rsvrd:16;       /* 2: reserved */
33     unsigned off1:8;         /* 1: offset in unit of fsblksize */
34     uint32_t off2;          /* 4: offset in unit of fsblksize */
35     unsigned len:24;         /* 3: length in unit of fsblksize */
36     unsigned addr1:8;        /* 1: address in unit of fsblksize */
37     uint32_t addr2;         /* 4: address in unit of fsblksize */
38 } xad_t;
39
40 #define MAXXLEN      ((1 << 24) - 1)
41
42 #define XTLOTSIZE    16
43 #define L2XTLOTSIZE  4
44
45 /* xad_t field construction */
46 #define XADoffset(xad, offset64)\
47 {\
48     (xad)->off1 = ((uint64_t)offset64) >> 32;\
49     (xad)->off2 = __cpu_to_le32((offset64) & 0xffffffff);\
50 }
51 #define XADaddress(xad, address64)\
52 {\
53     (xad)->addr1 = ((uint64_t)address64) >> 32;\
54     (xad)->addr2 = __cpu_to_le32((address64) & 0xffffffff);\
55 }
56 #define XADlength(xad, length32)      (xad)->len = __cpu_to_le24(length32)
57
58 /* xad_t field extraction */
59 #define offsetXAD(xad)\
60 ( ((int64_t)((xad)->off1)) << 32 | __le32_to_cpu((xad)->off2))
61 #define addressXAD(xad)\
62 ( ((int64_t)((xad)->addr1)) << 32 | __le32_to_cpu((xad)->addr2))
63 #define lengthXAD(xad)  __le24_to_cpu((xad)->len)
64
65 /* xad_t flags */
66 #define XAD_NEW      0x01    /* new */
67 #define XAD_EXTENDED 0x02    /* extended */
68 #define XAD_COMPRESSED 0x04 /* compressed with recorded length */
69 #define XAD_NOTRECORDED 0x08 /* allocated but not recorded */
70 #define XAD_COW     0x10    /* copy-on-write */
71
72
73 /* possible values for maxentry */
74 #define XTROOTINITSLOT_DIR  6
75 #define XTROOTINITSLOT     10
76 #define XTROOTMAXSLOT      18
77 #define XTPAGEMAXSLOT     256
78 #define XTENTRYSTART      2
79
80 /*
81 *   xtree page:
82 */
83 typedef union {
84     struct xthead {
85         int64_t next;          /* 8: */
86         int64_t prev;         /* 8: */
87
88         uint8_t flag;         /* 1: */
89         uint8_t rsvrd1;       /* 1: */
90         int16_t nextindex;     /* 2: next index = number of entries */

```

```
91             int16_t maxentry;          /* 2: max number of entries */
92             int16_t rsvrd2;           /* 2: */
93
94             pxd_t   self;              /* 8: self */
95     } header;                          /* (32) */
96
97     xad_t xad[XTPAGEMAXSLOT];         /* 16 * maxentry: xad array */
98 } xtpage_t;
99
100 #endif                                /* !_H_JFS_XTREE */
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef __H_DEBUG__
19
20 #define __H_DEBUG__
21
22 #include <stdio.h>
23
24 #ifdef TRACE
25 #define DBG_TRACE(a) {printf a; fflush(stdout);}
26 #else
27 #define DBG_TRACE(a) ;
28 #endif
29
30 #ifdef TRACE_IO
31 #define DBG_IO(a) {printf a; fflush(stderr);}
32 #else
33 #define DBG_IO(a) ;
34 #endif
35
36 #ifdef TRACE_ERROR
37 #define DBG_ERROR(a) {printf a; fflush(stdout);}
38 #else
39 #define DBG_ERROR(a) ;
40 #endif
41
42 #endif
```



```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #include "jfs_types.h"
19 #include "jfs_dmap.h"
20 #include "diskmap.h"
21
22 /*
23 *   budtab[]
24 *
25 *   used to determine the maximum free string in a character
26 *   of a wmap word. the actual bit values of the character
27 *   serve as the index into this array and the value of the
28 *   array at that index is the max free string.
29 *
30 */
31 static int8_t budtab[256] = {
32     3, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2,
33     2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
34     2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
35     2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
36     2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
37     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
38     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
39     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
40     2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
41     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
42     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
43     2, 1, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
44     2, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
45     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
46     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0,
47     2, 1, 1, 1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, -1
48 };
49
50 /*
51 *   NAME: ujfs_maxbuddy
52 *
53 *   FUNCTION: Determines the maximum string of free blocks within a word of the
54 *   wmap or pmap consistent with binary buddy.
55 *
56 *   PRE CONDITIONS:
57 *
58 *   POST CONDITIONS:
59 *
60 *   PARAMETERS:
61 *       cp           - Pointer to wmap or pmap word.
62 *
63 *   NOTES:
64 *
65 *   DATA STRUCTURES:
66 *
67 *   RETURNS: Maximum string of free blocks within word.
68 */
69 int8_t ujfs_maxbuddy(unsigned char *cp)
70 {
71     /*
72     *   Check if the wmap or pmap word is all free. If so, the free buddy size is
73     *   BUDMIN.
74     */
75     if (*((uint32_t *) cp) == 0) {
76         return (BUDMIN);
77     }
78
79     /*
80     *   Check if the wmap or pmap word is half free. If so, the free buddy size
81     *   is BUDMIN-1.
82     */
83     if (*((uint16_t *) cp) == 0 || *((uint16_t *) cp + 1) == 0) {
84         return (BUDMIN - 1);
85     }
86
87     /*
88     *   Not all free or half free. Determine the free buddy size through table
89     *   lookup using quarters of the wmap or pmap word.
90     */

```

```

91     return (MAX(MAX(budtab[*cp], budtab[*cp + 1]),
92                MAX(budtab[*cp + 2], budtab[*cp + 3])));
93 }
94
95 /*
96  * NAME: ujfs_adjtree
97  *
98  * FUNCTION: Calculate the tree of a dmap or dmapctl.
99  *
100 * PRE CONDITIONS:
101 *
102 * POST CONDITIONS:
103 *
104 * PARAMETERS:
105 *     cp        - Pointer to the top of the tree
106 *     l2leaves - Number of leaf nodes as a power of 2
107 *     l2min    - Number of disk blocks actually covered by a leaf of the tree;
108 *               specified as a power of 2
109 *
110 * NOTES: This routine first works against the leaves of the tree to calculate
111 *        the maximum free string for leaf buddies. Once this is accomplished the
112 *        values of the leaf nodes are bubbled up the tree.
113 *
114 * DATA STRUCTURES:
115 *
116 * RETURNS:
117 */
118 int8_t ujfs_adjtree(int8_t * treep, int32_t l2leaves, int32_t l2min)
119 {
120     int32_t nleaves, leaf_index, l2max, nextb, bsize, index;
121     int32_t l2free, leaf, num_this_level, nextp;
122     int8_t *cp0, *cp1, *cp = treep;
123
124     /*
125      * Determine the number of leaves of the tree and the
126      * index of the first leaf.
127      * Note: I don't know why the leaf_index calculation works, but it does.
128      */
129     nleaves = (1 << l2leaves);
130     leaf_index = (nleaves - 1) / 3;
131
132     /*
133      * Determine the maximum free string possible for the leaves.
134      */
135     l2max = l2min + l2leaves;
136
137     /*
138      * Try to combine buddies starting with a buddy size of 1 (i.e. two leaves).
139      * At a buddy size of 1 two buddy leaves can be combined if both buddies
140      * have a maximum free of l2min; the combination will result in the
141      * left-most buddy leaf having a maximum free of l2min+1. After processing
142      * all buddies for a certain size, process buddies at the next higher buddy
143      * size (i.e. current size * 2) and the next maximum free
144      * (current free + 1). This continues until the maximum possible buddy
145      * combination yields maximum free.
146      */
147     for (l2free = l2min, bsize = 1; l2free < l2max; l2free++, bsize = nextb) {
148         nextb = bsize << 1;
149
150         for (cp0 = cp + leaf_index, index = 0; index < nleaves;
151              index += nextb, cp0 += nextb) {
152             if (*cp0 == l2free && *(cp0 + bsize) == l2free) {
153                 *cp0 = l2free + 1;
154                 *(cp0 + bsize) = -1;
155             }
156         }
157     }
158
159     /*
160      * With the leaves reflecting maximum free values bubble this information up
161      * the tree. Starting at the leaf node level, the four nodes described by
162      * the higher level parent node are compared for a maximum free and this
163      * maximum becomes the value of the parent node. All lower level nodes are
164      * processed in this fashion then we move up to the next level (parent
165      * becomes a lower level node) and continue the process for that level.
166      */
167     for (leaf = leaf_index, num_this_level = nleaves >> 2; num_this_level > 0;
168          num_this_level >>= 2, leaf = nextp) {
169         nextp = (leaf - 1) >> 2;
170
171         /*
172          * Process all lower level nodes at this level setting the value of the
173          * parent node as the maximum of the four lower level nodes.
174          */
175         for (cp0 = cp + leaf, cp1 = cp + nextp, index = 0;
176              index < num_this_level; index++, cp0 += 4, cp1++) {
177             *cp1 = TREEMAX(cp0);
178         }
179     }
180

```

```

181     return (*cp);
182 }
183
184 /*
185  * NAME: ujfs_complete_dmap
186  *
187  * FUNCTION: Fill in rest of dmap fields from wmap/pmap already initialized.
188  *
189  * PARAMETERS:
190  *     dmap_page     - dmap to complete
191  *     blkno         - starting block number for this dmap
192  *     treemax       - will be filled in with max free for this dmap
193  *
194  * RETURNS: NONE
195  */
196 void ujfs_complete_dmap(struct dmap *dmap_page, int64_t blkno, int8_t *treemax)
197 {
198     struct dmaptree *tp;
199     int8_t *cp;
200     int32_t index;
201
202     dmap_page->start = blkno;
203
204     tp = &dmap_page->tree;
205     tp->height = 4;
206     tp->leafidx = LEAFIND;
207     tp->nleaves = LPERDMAP;
208     tp->l2nleaves = L2LPERDMAP;
209     tp->budmin = BUDMIN;
210
211     /*
212      * Pick up the pointer to the first leaf of the dmap tree.
213      */
214     cp = tp->stree + tp->leafidx;
215
216     /*
217      * Set the initial state for the leaves of the dmap tree. They will reflect
218      * the current allocation state of the wmap words.
219      */
220     for (index = 0; index < LPERDMAP; index++) {
221         *(cp + index) = ujfs_maxbuddy((unsigned char *) &dmap_page->wmap[index]);
222     }
223
224     /*
225      * With the leaves of the dmap initialized adjust (initialize) the dmap's
226      * tree.
227      */
228     *treemax = ujfs_adjtree(tp->stree, L2LPERDMAP, BUDMIN);
229 }
230
231 /*
232  * NAME: ujfs_idmap_page
233  *
234  * FUNCTION: Initialize one dmap page
235  *
236  * POST CONDITIONS: Blocks which don't actually exist in the aggregate will be
237  * marked as allocated in the dmap page. The total number of blocks will
238  * only account for the actually existing blocks.
239  *
240  * PARAMETERS:
241  *     map_page     - pointer to page of map
242  *     nblocks      - number of blocks this page
243  *
244  * RETURNS: NONE
245  */
246 void ujfs_idmap_page(struct dmap *map_page, uint32_t nblocks)
247 {
248     uint32_t index, nwords, bit;
249
250     map_page->nblocks = map_page->nfree = nblocks;
251
252     /*
253      * Partial dmap page?
254      * If there are not enough blocks to cover an entire dmap page the ones
255      * which represent blocks which don't exist will be marked as allocated.
256      *
257      * nwords will indicate the first word beyond the end of existing blocks
258      * bit will indicate if this block does not fall on a 32-bit boundary
259      */
260     nwords = nblocks / DBWORD;
261     bit = nblocks % DBWORD;
262
263     if (bit) {
264         /*
265          * Need to mark a partial word allocated
266          */
267         map_page->wmap[nwords] = map_page->pmap[nwords] = ONES >> bit;
268         nwords++;
269     }
270

```

```
271     /*
272     * Set the rest of the words in the page to ONES.
273     */
274     for (index = nwords; index < LPERDMAP; index++) {
275         map_page->pmap[index] = map_page->wmap[index] = ONES;
276     }
277 }
278
279 /*
280 * NAME: ujfs_getagl2size
281 *
282 * FUNCTION: Determine log2(allocation group size) based on size of aggregate
283 *
284 * PARAMETERS:
285 *     size - Number of blocks in aggregate
286 *     aggr_block_size - Aggregate block size
287 *
288 * RETURNS: log2(allocation group size) in aggregate blocks
289 */
290 int32_t ujfs_getagl2size(int64_t size, int32_t aggr_block_size)
291 {
292     int64_t sz;
293     int64_t m;
294     int32_t l2sz;
295
296     if (size < BPERDMAP * MAXAG) {
297         return (L2BPERDMAP);
298     }
299
300     m = ((uint64_t) 1 << (64 - 1));
301     for (l2sz = 64; l2sz >= 0; l2sz--, m >>= 1) {
302         if (m & size) {
303             break;
304         }
305     }
306
307     sz = (int64_t) 1 << l2sz;
308     if (sz < size) {
309         l2sz += 1;
310     }
311
312     return (l2sz - L2MAXAG);
313 }
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_DISKMAP
19 #define H_DISKMAP
20
21 #include <jfs_types.h>          /* int8_t, int32_t, int64_t, uint32_t */
22
23 int8_t ujfs_maxbuddy(unsigned char *);
24 int8_t ujfs_adjtree(int8_t *, int32_t, int32_t);
25 void ujfs_complete_dmap(struct dmap *, int64_t, int8_t *);
26 void ujfs_idmap_page(struct dmap *, uint32_t);
27 int32_t ujfs_gettagl2size(int64_t, int32_t);
28
29 #endif                          /* H_DISKMAP */
```



```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_JFS_FSCKCBBL
19 #define _H_JFS_FSCKCBBL
20
21 /*
22 *   This structure resides in the first page (aka the block map control page)
23 *   of the fsck in-aggregate workspace.  JFS Clear Bad Block List utility
24 *   processing writes to this record, then fsck reads it and reports to the
25 *   caller.
26 */
27 struct fsckcbbbl_record {
28     char eyecatcher[8];
29     char avail_1[4];          /* 4 */
30     int32_t cbbbl_retcode;   /* JFS Clear Bad Block List utility
31                               * return code
32                               */
33     int32_t fs_blksize;      /* aggregate block size */
34     int32_t lv_blksize;      /* device block size */
35     int32_t fs_lv_ratio;     /* fs_blksize/lv_blksize */
36     int64_t fs_last_metablk; /*
37                               * last fs block we won't try to relocate
38                               * because it holds fixed-location metadata
39                               */
40     int64_t fs_first_wspblk; /*
41                               * first fs block we won't try to relocate
42                               * because it holds fsck workspace or the
43                               * inline journal log
44                               */
45     int32_t total_bad_blocks; /* count of bad blocks in LVM's list
46                               * at beginning of Bad Block List utility
47                               * processing
48                               */
49     int32_t resolved_blocks; /* count of bad blocks:
50                               * - for which the data has been relocated,
51                               * - which are now allocated to the bad block
52                               *   inode, and
53                               * - which the LVM has been told to forget
54                               */
55     int32_t reloc_extents;   /* count of relocated extents */
56     int64_t reloc_blocks;    /* count of blocks in relocated extents */
57     int32_t LVM_lists;       /* count of bad block lists maintained by LVM
58                               * according to the last query
59                               */
60     char bufptr_eyecatcher[8];
61     void *clrbblks_agg_recptr; /* addr of clrbblks aggregate record */
62     void *imapInoPtr;         /* addr of imap inode buffer */
63     void *imapCtlPtr;        /* addr of imap control page buffer */
64     void *imapLeafPtr;       /* addr of imap leaf page buffer */
65     void *iagPtr;            /* addr of iag buffer */
66     void *inoExtPtr;          /* addr of inode extent buffer */
67     char avail_2[28];        /* 28 */
68 };                          /* total = 128 bytes */
69
70 #endif                       /* _H_JFS_FSCKCBBL */

```



```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_FSCKWSP
19 #define H_FSCKWSP
20
21 #include "fscklog.h"
22 #include <jfs_dmap.h>
23 #include <jfs_dtree.h>
24 #include <jfs_xtree.h>
25 #include <jfs_filsys.h>
26 #include <jfs_imap.h>
27 #include <jfs_dinode.h>
28 #include "fsck_base.h"
29 #include "fsckcbl.h"
30
31 /*-----
32 *   fsck uses the structures and types in this file to define its workspace.
33 *
34 *-----
35 */
36
37 /*-----
38 *   type dtree_Q_elem describes one node which is to be visited in the
39 *   traversal of a directory B+ Tree
40 *
41 *-----
42 */
43 struct dtreeQelem {
44     struct dtreeQelem *next;
45     struct dtreeQelem *prev;
46     int64_t node_addr;
47     pxd_t node_pxd;
48     uint8_t node_key_len;
49     UniChar node_key[JFS_NAME_MAX];
50     int8_t node_level;
51     uint32_t node_size;
52 };
53
54 #define dtreeQ_elem_length sizeof(struct dtreeQelem)
55
56 /*-----
57 *   type tree_Q_elem describes one node which is to be visited in the
58 *   traversal of a non-directory B+ Tree
59 *
60 *-----
61 */
62 struct treeQelem {
63     struct treeQelem *next;
64     struct treeQelem *prev;
65     int64_t node_addr;
66     int64_t node_first_offset;
67     pxd_t node_pxd;
68     int8_t node_level;
69     int8_t last_ext_uneven;
70 };
71
72 #define treeQ_elem_length sizeof(struct treeQelem)
73
74 /*-----
75 *   type dupall_blkrec describes one aggregate block which has been found to
76 *   have more than 1 inode claiming to own it.
77 *
78 *-----
79 */
80 struct dupall_blkrec {
81     int64_t blk_number;
82     uint8_t first_ref_resolved;
83     uint8_t avail_byte;
84     uint16_t owner_count;
85     struct dupall_blkrec *next;
86 };
87
88 #define dupall_blkrec_length sizeof(struct dupall_blkrec)
89
90 /*-----
91 *   Each (4096 byte) page in the block map describes 32768

```

```

91  * blocks in the aggregate.
92  */
93
94  /*
95  * the fsck Workspace Block Map control page
96  *
97  */
98
99  struct fsck_blk_map_hdr {
100     /*
101     * the 1st 1024 bytes are the clrbblks to fsck
102     * communication area
103     */
104     struct fsckcbbbl_record cbbblrec; /* 128 */
105     char avail_0[896]; /* 896 */
106     /* 1024 byte boundary */
107     /*
108     * the next 2048 bytes are the fsck block map
109     * header area
110     */
111     struct {
112         char eyecatcher[8]; /* 8 */
113         int32_t last_entry_pos; /* 4 */
114         int32_t next_entry_pos; /* 4 */
115         char start_time[20]; /* 20 */
116         char end_time[20]; /* 20 */
117         int32_t return_code; /* 4 */
118         char *super_buff_addr; /* 4 -- for the dump reader */
119         char *agg_record_addr; /* 4 -- for the dump reader */
120         char *bmap_record_addr; /* 4 -- for the dump reader */
121         char avail_1[8]; /* 8 */
122         int64_t fscklog_agg_offset; /* 8 */
123         int8_t fscklog_full; /* 1 */
124         int8_t fscklog_buf_allocated; /* 1 */
125         int8_t fscklog_buf_alloc_err; /* 1 */
126         char avail_2[1]; /* 1 */
127         int32_t num_logwrite_errors; /* 4 */
128         char avail_3[32]; /* 32 */
129         struct fscklog_error logerr[125]; /* 1920 = 120 * 16 */
130     } hdr; /* subtotal, 2048 bytes */
131     /*
132     * the last 1024 bytes are available
133     */
134     char avail_5[1024]; /* 1024 */
135 }; /* total: 4096 bytes */
136
137 #define fbmh_eyecatcher_string "wspblkmp"
138
139 /*
140 * a page in the fsck Workspace Block Map
141 */
142 struct fsck_blk_map_page {
143     uint32_t fsck_blkmap_words[1024]; /* 4096 bytes */
144 };
145
146 /*-----
147 * structure blkmap_wspace describes the portion of the aggregate
148 * record allocated for use when verifying
149 * the aggregate block map.
150 *
151 */
152 struct blkmap_wspace {
153     uint32_t dmap_map[LPERDMAP]; /* bit map of blk allocations */
154     int8_t dmap_wsp_tree[TREESIZE - LPERDMAP];
155     int8_t dmap_wsp_leafs[LPERDMAP];
156     int8_t L0_wsp_tree[CTLTREESIZE - LPERCTL];
157     int8_t L0_wsp_leafs[LPERCTL];
158     int8_t L1_wsp_tree[CTLTREESIZE - LPERCTL];
159     int8_t L1_wsp_leafs[LPERCTL];
160     int8_t L2_wsp_tree[CTLTREESIZE - LPERCTL];
161     int8_t L2_wsp_leafs[LPERCTL];
162     int64_t AG_free[MAXAG];
163 };
164
165 /*-----
166 * structure fsck_iag_record describes one Inode Allocation
167 * Group belonging to an inode table in
168 * the aggregate
169 *
170 */
171 struct fsck_iag_record {
172     uint32_t AG_num;
173     uint32_t backed_inodes;
174     uint32_t unused_backed_inodes;
175     uint32_t unbacked_extents;
176 };
177
178 #define iag_record_length sizeof(struct fsck_iag_record)
179
180 /*-----

```

```

181 * structure fsck_ag_record describes one Allocation Group
182 *
183 *           as it relates to an inode table in the
184 *           aggregate
185 */
186 struct fsck_ag_record {
187     uint32_t backed_inodes;
188     uint32_t unused_backed_inodes;
189     int32_t frext_list_first;
190     int32_t frext_list_last;
191     int32_t frext_list_len;
192     int32_t frino_list_first;
193     int32_t frino_list_last;
194     int32_t frino_list_len;
195     int8_t frext_list_bad;
196     int8_t frino_list_bad;
197     char unused[14];
198 };
199
200 #define ag_record_length sizeof(struct fsck_ag_record)
201
202 /*-----
203 * structure fsck_iam_record describes one Inode Allocation Map
204 *           in the aggregate
205 *
206 */
207 struct fsck_iam_record {
208     int64_t first_leaf_offset; /* offset of first leaf for imap inode */
209     int8_t imap_is_rootleaf;
210     int8_t friag_list_bad;
211     char unused[2];
212     int32_t bkd_inodes; /* count of backed inodes */
213     int32_t unused_bkd_inodes; /* count of available backed inodes */
214     int32_t num_iags; /* number of IAGs in the inode tbl */
215     int32_t friag_list_last;
216     int32_t friag_list_len;
217     struct fsck_iag_record *iag_tbl; /* ptr to info by iag */
218     struct fsck_ag_record *ag_tbl; /* ptr to info by ag */
219 };
220
221 #define iam_record_length sizeof(struct fsck_iam_record)
222
223 /*-----
224 * structure fsck_inode_record describes one inode belonging to
225 *           the aggregate or to a fileset
226 *           in the aggregate
227 *
228 */
229
230 /*
231 * the possible values for inode_type
232 * (since inode_type is 4 bits wide, valid range is 0-15)
233 */
234 #define unused_inode 0
235 #define file_inode 1
236 #define strm_descr_inode 2
237 #define stream_inode 3
238 #define directory_inode 4
239 #define symlink_inode 5
240 #define block_special_inode 6
241 #define char_special_inode 7
242 #define FIFO_inode 8
243 #define SOCK_inode 9
244 #define metadata_inode 14
245 #define unrecognized_inode 15
246
247 #define add_direntiry_extension 'A'
248 #define rmv_direntiry_extension 'R'
249 #define rmv_badentry_extension 'B'
250 #define parent_extension 'P'
251
252 struct fsck_inode_ext_record {
253     char ext_type; /* 1 : extends type (see constants above) */
254     unsigned ino_type:4; /* .5: see constants above
255 * describes the inonum inode.
256 * This is only interesting when
257 * type is rmv_direntiry_extension.
258 */
259     unsigned unused_byte:20; /* 2.5: unused */
260     struct fsck_inode_ext_record *next; /*
261 * 4 : addr next extension record
262 */
263     uint32_t inonum; /* 4 : inode number */
264 };
265
266 /*
267 * the inode record data
268 */
269 struct fsck_inode_record {
270     uint32_t inonum; /* 4.0 : key */

```

```

271     /* 3 : 24 flag bits */
272     unsigned in_use:1;
273     unsigned selected_to_rls:1;
274     unsigned crct_link_count:1;
275     unsigned crct_prnt_inonum:1;
276
277     unsigned adj_entries:1;
278     unsigned cant_chkea:1;
279     unsigned clr_ea_fld:1;
280     unsigned clr_acl_fld:1;
281
282     unsigned ignore_ea_blks:1;
283     unsigned ignore_acl_blks:1;
284     unsigned ignore_alloc_blks:1;
285     unsigned inline_data_err:1;
286
287     unsigned reconnect:1;
288     unsigned unxpctd_prnts:1;
289     unsigned badblk_inode:1;
290     unsigned involved_in_dups:1;
291
292     unsigned crct_cumm_blks:1;
293     unsigned avail_flagbit_1:1;
294     unsigned inlineea_on:1;
295     unsigned inlineea_off:1;
296
297     unsigned avail_flagbits:4;
298     /*      end of flag bits      */
299     int8_t inode_type; /* 1 : see constants above      */
300
301     int32_t link_count; /* 4 : Stored link count is added to this.
302                        *      Decrementd for each observed link.
303                        */
304     uint32_t parent_inonum; /* 4 : If this is a directory inode and
305                        *      the unxpctd_prnts bit is set, then
306                        *      this field contains the inode number
307                        *      stored in the parent field on disk.
308                        *      Otherwise, this is the inode number
309                        *      of the first observed parent.
310                        */
311     int64_t cumm_blocks; /* 8 : number of blocks allocated to the
312                        *      inode and, if this is a directory, to each
313                        *      inode described in an entry
314                        */
315     int32_t dtree_level; /* 4 : distance, in the fileset directory tree,
316                        *      from the root (inode 2)
317                        */
318     struct fsck_inode_ext_record *ext_rec; /* 4.0: extension record address */
319 };
320
321 #define inode_record_length sizeof(struct fsck_inode_record)
322
323 #define inode_is_metadata( X ) \
324     ( (int8_t) ((X)->inode_type) == (int8_t) metadata_inode )
325
326 /*-----
327  * type inode_tbl_t is an array of pointers to inode records
328  *
329  * There is one of these tables for each allocated inode extent which
330  * has at least one inode in use.
331  */
332
333 struct inode_tbl_t {
334     char eyecatcher[8]; /* 8.0 */
335     struct fsck_inode_record *inorectbl[32]; /* 128 */
336 };
337
338 #define inode_tbl_length sizeof( struct inode_tbl_t )
339
340 /*-----
341  * type inode_ext_tbl_t is an array of pointers to inode tables
342  *
343  * There is one of these tables for each allocated IAG which
344  * has at least one allocated extent with an inode in use.
345  */
346
347 struct inode_ext_tbl_t {
348     char eyecatcher[8]; /* 8.0 */
349     struct inode_tbl_t *inotbl[128]; /* 512 */
350 };
351
352 #define inode_ext_tbl_length sizeof( struct inode_ext_tbl_t )
353
354 /*-----
355  * type IAG_tbl_t is an array of pointers to inode extent tables
356  *
357  * There is one of these tables for each fileset in the aggregate,
358  * and is allocated dynamically when the number of IAGs in the
359  * fileset Inode Map is known.
360  */

```

```

361  */
362  struct IAG_tbl_t {
363      char eyecatcher[8]; /* 8.0 */
364      struct inode_ext_tbl_t *inoext_tbl[1]; /* 4.0 */
365  };
366
367  /*-----
368  * type wsp_ext_rec describes one extent of storage which fsck
369  * has allocated for its workspace.
370  *
371  */
372  struct wsp_ext_rec {
373      struct wsp_ext_rec *next;
374      int8_t from_high_memory;
375      int8_t for_logredo;
376      char avail[2];
377      uint32_t extent_length;
378      char *extent_addr;
379      uint32_t last_byte_used;
380  };
381
382  /*-----
383  * type recon_buf_record describes one dnode I/O buffer used during
384  * inode reconnect processing
385  *
386  */
387  struct recon_buf_record {
388      dtpage_t dnode_buf;
389      struct recon_buf_record *stack_next;
390      int64_t dnode_blkoff;
391      int64_t dnode_byteoff;
392      char reserved[12];
393  };
394
395  /*-----
396  * type treeStack_record describes one stack element used during DASD
397  * limits processing
398  *
399  */
400  struct treeStack_record {
401      struct fsck_inode_record *inorec;
402      struct treeStack_record *next;
403  };
404
405  #define treeStack_elem_length sizeof(struct treeStack_record)
406
407  /*-----
408  * structure fsck_agg_record describes the aggregate
409  */
410
411  struct fsck_agg_record {
412      char eyecatcher[8];
413      uint32_t ondev_jlog_fsblk_length;
414      int64_t ondev_jlog_fsblk_offset;
415      int64_t ondev_wsp_byte_length;
416      int64_t ondev_wsp_byte_offset;
417      uint32_t ondev_wsp_fsblk_length;
418      int64_t ondev_wsp_fsblk_offset;
419      uint32_t ondev_fscklog_fsblk_length;
420      int64_t ondev_fscklog_fsblk_offset;
421      int64_t ondev_fscklog_byte_length;
422      int64_t ondev_fscklog_byte_offset;
423      uint32_t log2_blksize;
424      uint32_t blkspcrpg;
425      uint32_t log2_blkspcrpg;
426      uint32_t log2_blkspcrag;
427      int64_t sb_agg_fsblk_length;
428      uint32_t inode_stamp; /* taken from the agg self inode, every
429                          * in-use, valid inode must have a
430                          * matching di_inostamp.
431                          */
432      int inode_count; /* num inodes owned by the aggregate */
433      int inodes_inuse; /* num aggregate nodes now in use */
434      int fset_inodes_inuse; /* num fileset inodes in use */
435      int fset_inodes_seen; /* num fileset inodes scanned. */
436      int fset_inode_count; /* num inodes in the filesets */
437      int64_t dup_block_count; /* num multiply-allocated
438                          * blocks seen
439                          */
440      int64_t unresolved_lstref_count; /* num unresolved first
441                          * references to multiply allocated
442                          * blocks
443                          */
444      char aggreg_rsvd0[4];
445      int64_t blocks_for_inodes; /* by fsck's count */
446      int64_t blocks_for_files; /* by fsck's count */
447      int64_t blocks_for_dirs; /* by fsck's count */
448      int64_t blocks_for_eas; /* by fsck's count */
449      int64_t blocks_for_acls; /* by fsck's count */
450      uint64_t inodes_in_aggregate; /* by fsck's count */

```

```

451     uint64_t files_in_aggregate; /* by fsck's count */
452     uint64_t dirs_in_aggregate; /* by fsck's count */
453     int64_t free_blocks_in_aggregate; /* by fsck's count */
454     int64_t blocks_used_in_aggregate; /* by fsck's count */
455     uint64_t blocks_this_fset; /* by fsck's count */
456
457     int32_t logredo_rc; /* logredo return code */
458     int32_t logformat_rc; /* logformat return code */
459     char aggrec_rsvd2[8];
460
461     struct { /* data for the current inode */
462         char eyecatcher[8];
463         int64_t all_blks; /* all blocks allocated */
464         int64_t data_blks; /* blocks allocated to data */
465         uint64_t data_size; /* byte offset of last data if file */
466         int64_t ea_blks; /* blocks allocated to EAs */
467         int64_t acl_blks; /* blocks allocated to ACLs */
468         uint16_t ea_inline;
469         uint16_t inline_ea_offset;
470         uint16_t inline_ea_length;
471         uint16_t acl_inline;
472         uint16_t inline_acl_offset;
473         uint16_t inline_acl_length;
474         uint16_t in_inode_data_length; /* length of inline data or of
475                                         * xad list in the inode
476                                         */
477         char aggrec_rsvd4[2];
478     } this_inode;
479
480     struct wsp_ext_rec *recon_buf_extent;
481     struct recon_buf_record *recon_buf_stack;
482
483     char agg_imap_eyecatcher[8];
484     struct fsck_iam_record agg_imap; /* describes the aggregate imap */
485
486     char fset_imap_eyecatcher[8];
487     struct fsck_iam_record fset_imap; /* describes the fileset imap */
488
489     uint32_t num_ag; /* number of AGs in the aggregate */
490     int ag_blk_size; /* aggregate block size */
491
492     short int aggrec_rsvd6[1];
493     char delim_char; /* path delimiter char */
494     /* unicharacter equivalents */
495     UniChar *UniChar_lsfname; /* equiv to lost+found */
496     UniChar *UniChar_LSFNAME; /* equiv to LOST+FOUND */
497     /* end of unicharacter equivalents */
498     int64_t lsfname_dasdused_adjustment;
499     uint32_t lsfname_inonum; /* inode number of /lost+found/ */
500     uint32_t avail_inonum; /* inode number of an inode
501                             * observed to be allocated and not
502                             * in use (if any)
503                             */
504     short int parm_options[16]; /* parms specified */
505     char aggrec_rsvd7[2]; /* pad to quadword boundary */
506     char effective_msg_level;
507     char aggrec_rsvd8[7]; /* pad to quadword boundary */
508     char flags_eyecatcher[8];
509     /* 4 bytes of flags */
510     unsigned parm_options_nologredo:1; /* first byte, first nibble */
511     unsigned processing_readonly:1;
512     unsigned processing_readwrite:1;
513     unsigned messaging_verbose:1;
514     unsigned superbk_ok:1; /* first byte, second nibble */
515     unsigned aggregate_is_mounted:1;
516     unsigned ag_modified:1;
517     unsigned ag_dirty:1;
518     unsigned lsfname_ok:1; /* second byte, first nibble */
519     unsigned cant_write_primary_sb:1;
520     unsigned cant_write_secondary_sb:1;
521     unsigned cant_write_primary_ait:1;
522     unsigned nonmeta_write_error_msg:1; /* second byte, second nibble */
523     unsigned fsck_is_done:1;
524     unsigned device_is_open:1;
525     unsigned device_is_locked:1;
526     unsigned primary_ait_4part1:1; /* third byte, first nibble */
527     unsigned primary_ait_4part2:1;
528     unsigned ait_aim_update_failed:1;
529     unsigned corrections_needed:1;
530     unsigned corrections_approved:1; /* third byte, second nibble */
531     unsigned avail_inode_found:1;
532     unsigned initializing_fscklog:1;
533     unsigned warning_pending:1;
534     unsigned high_mem_allocated:1; /* fourth byte, first nibble */
535     unsigned rootdir_rebuilt:1;
536     unsigned active_dasd_limits:1;
537     unsigned stdout_redirected:1;
538     unsigned prime_dasd_limits:1; /* fourth byte, second nibble */
539     unsigned unused_4_2_3; /*
540                             */
541     /* end of flag bytes */

```

```

541     int path_buffer_length; /* length of path_buffer */
542     char *path_buffer;     /* storage allocated (if any) for
543                            * assembling the string containing
544                            * an inode's path for a message
545                            */
546     xtpage_t *prim_nodeptr; /* storage allocated (if any) for
547                            * an xtpage of an inode in the Primary Agg
548                            * Inode table
549                            */
550     xtpage_t *second_nodeptr; /* storage allocated (if any) for
551                              * an xtpage of an inode in the Secondary
552                              * Agg Inode table
553                              */
554     int64_t lowest_valid_fset_datablk; /* the lowest (aggregate
555                                        * blocksized) block in the aggregate
556                                        * AFTER the blocks used for initial
557                                        * metadata
558                                        */
559     int64_t highest_valid_fset_datablk; /* the highest (aggregate
560                                        * blocksized) block in the aggregate before
561                                        * the beginning of the fsck workspace
562                                        */
563     struct fsck_inode_ext_record *free_inode_extens; /* available workspace inode record
564                                                       * extensions.
565                                                       */
566     struct fsck_inode_ext_record *inode_reconn_extens; /* list of inode record
567                                                         * extensions describing the directory
568                                                         * entries to be added when reconnecting
569                                                         * inodes
570                                                         */
571     struct treeStack_record *treeStack; /* head of the stack used during tree
572                                         * traversal
573                                         */
574     struct treeStack_record *free_treeStack; /* available tree stack elements */
575     int32_t tree_height;
576     struct dtreeQelem *dtreeQ_front;
577     struct dtreeQelem *dtreeQ_back;
578     struct dtreeQelem *free_dtreeQ; /* available dtree stack elements */
579     struct treeQelem *treeQ_front;
580     struct treeQelem *treeQ_back;
581     struct treeQelem *free_treeQ; /* available tree stack elements */
582     struct dupall_blkrec *dup_alloc_lst; /* list of multiply allocated blocks */
583     struct dupall_blkrec *free_dupall_blkrec; /* free dupall_blkrec records */
584     struct wsp_ext_rec *wsp_extent_list; /* linked list of records describing
585                                         * the workspace extents.
586                                         */
587     char aggrec_rsvda[4]; /* pad to quadword boundary */
588
589     char AIT_eyecatcher[8]; /*
590
591     struct inode_tbl_t *AIT_ext0_tbl; /* ptr to table for Agg Inode Extent 0 */
592     int32_t agg_last_inoidx; /* used for find first, find next */
593
594     char FSIT_eyecatcher[8];
595     struct IAG_tbl_t *FSIT_IAG_tbl; /* ptr to table for FSet IAGs */
596     int32_t fs_last_iagidx; /* used for find first, find next */
597     int32_t fs_last_extidx; /* used for find first, find next */
598     int32_t fs_last_inoidx; /* used for find first, find next */
599
600     struct {
601         /* data for accessing all allocated inodes in an
602          * inode table sequentially
603          */
604         char eyecatcher[8];
605         int32_t this_iagnum; /* ordinal of the current iag */
606         uint32_t this_inoidx; /* ordinal of the current inode */
607         uint32_t rootleaf_imap; /* 0 => not a rootleaf imap */
608         xtpage_t *this_mapleaf; /* current mapleaf in buffer */
609         int32_t iagidx_now; /* index into current mapleaf */
610         int32_t iagidx_max;
611         struct iag *iagptr; /* current iag in buffer */
612         uint32_t extidx_now; /* index into current iag */
613         uint32_t extidx_max;
614         struct dinode *extptr; /* current inode extent in buffer */
615         uint32_t inoidx_now; /* index into current inode extent */
616         uint32_t inoidx_max;
617         struct dinode *inoptr; /* current inode in buffer */
618     } fais; /* (For Allocated Inodes Sequentially) */
619
620     char aggrec_rsvde[4]; /* pad to quadword boundary */
621
622     /*
623     * The fsck I/O buffer information is below. Certain information
624     * is kept for each of the buffers:
625     *
626     * %_buf_ptr = buffer address
627     * %_buf_length = buffer length
628     * %_buf_data_len = length of data read into buffer at last read
629     * %_agg_offset = byte offset in aggregate of buffer contents
630     * %_buf_write = flag: !0 => buffer contents modified since
631                   last write to device
632     */

```

```

631     * where % = {blkmp, iag, mapleaf, ino, node, wsp}
632     *
633     *
634     * N.B. To ensure proper boundary alignment, each IO buffersize
635     *      must be an even number of pages (i.e. 4096 byte pages).
636     *
637     */
638
639     /* ***** very large, multi purpose BUFFER ***** */
640
641     #define VLARGE_BUFSIZE (2 * MEMSEGSIZE)
642
643     #define NOT_CURRENTLY_USED 0
644     #define USED_FOR_EA_BUF 1 /* phase 1 late */
645     #define USED_FOR_DIRPAGE_BUFS 2 /* phase 6 */
646     #define USED_FOR_INOEXT_BUF 3 /* phase 1 early, phase 7 */
647     #define USED_FOR_SUPER_VALIDATION 4 /* initial processing */
648
649     char vlarge_info_eyecatcher[8];
650     char *vlarge_buf_ptr;
651     uint32_t vlarge_buf_length;
652     uint32_t vlarge_current_use;
653     char aggreg_rsvdf[4]; /* pad to quadword boundary */
654
655     /* ***** fsck log BUFFER ***** */
656
657     #define FSCKLOG_BUFSIZE (2 * BYTESPERPAGE)
658
659     char fscklog_info_eyecatcher[8];
660     struct fsck_blk_map_page *fscklog_buf_ptr;
661     uint32_t fscklog_buf_length;
662     uint32_t fscklog_buf_data_len;
663     int64_t fscklog_agg_offset;
664     int64_t fscklog_log_offset;
665     struct fscklog_entry_hdr *fscklog_last_msghdr;
666     int8_t fscklog_full;
667     int8_t fscklog_buf_allocated;
668     int8_t fscklog_buf_alloc_err;
669     char aggreg_rsvdg[5]; /* pad to quadword boundary */
670
671     /* ***** fsck block map BUFFER *****
672     *
673     * N.B. If the fsck workspace block map is instantiated in
674     *      dynamic storage, then the following describe the entire
675     *      block map. However, if the fsck workspace block map
676     *      is instantiated within the aggregate, then the following
677     *      describe the I/O buffer associated with it.
678     */
679
680     #define BLKMP_IO_BUFSIZE 4 * BYTESPERPAGE
681
682     char blkmp_info_eyecatcher[8];
683     struct fsck_blk_map_page *blkmp_buf_ptr;
684     uint32_t blkmp_buf_length;
685     uint32_t blkmp_buf_data_len;
686     int64_t blkmp_agg_offset;
687     int64_t blkmp_blkmp_offset;
688     int8_t blkmp_buf_write;
689     struct fsck_blk_map_hdr *blkmp_ctlptr;
690     int32_t blkmp_pagecount;
691
692     /* ***** BLOCK MAP DMAP PAGE I/O BUFFER ***** */
693     /*
694     * note: The EA buffer is treated differently from the
695     *      other I/O buffers allocated for fsck.
696     *      Specifically,
697     *
698     *      - it is larger than any other fsck I/O buffer
699     *        (sized to accomodate the largest legal ea)
700     *      - since it is exactly 1 memory segment, it is
701     *        not allocated as a normal wsp extent (described
702     *        by an fer on the wsp_extent_list) but is
703     *        allocated as a special case.
704     *      - this buffer is released as soon as phase 1 is
705     *        completed (does not persist until fsck has
706     *        finished processing)
707     *      - BTW the EA data is read by fsck, but never
708     *        written.
709     */
710
711     #define EA_IO_BUFSIZE 16*BYTESPERPAGE
712
713     char ea_info_eyecatcher[8];
714     char *ea_buf_ptr;
715     uint32_t ea_buf_length;
716     uint32_t ea_buf_data_len;
717     int64_t ea_agg_offset;
718
719     char aggreg_rsvdh[4]; /* pad to quadword boundary */
720

```



```

721     /* ***** IAG I/O BUFFER ***** */
722     /*
723     * note: This is not actually a unique buffer. Since
724     * the Inode Allocation Maps verification is
725     * completed before the Block Allocation Map
726     * verification, buffer space is allocated and
727     * used first for Inode Alloc Map verification
728     * processing and then for Block Allocation Map
729     * verification.
730     */
731
732     #define IAG_IO_BUFSIZE 4 * BYTESPERPAGE
733
734     char iag_info_eyecatcher[8];
735     char *iag_buf_ptr;
736     uint32_t iag_buf_length;
737     uint32_t iag_buf_data_len;
738     int64_t iag_agg_offset;
739     uint32_t iag_buf_1st_inode; /* inode number of 1st in buffer */
740     uint32_t iag_fsnum;
741     int8_t iag_for_aggregate;
742     int8_t iag_which_it;
743     int8_t iag_buf_write;
744     char aggrec_rsvdi[9]; /* pad to quadword boundary */
745
746     /* ***** MAP CONTROL PAGE I/O BUFFER ***** */
747
748     #define MAPCTL_IO_BUFSIZE BYTESPERPAGE
749
750     char mapctl_info_eyecatcher[8];
751     char *mapctl_buf_ptr;
752     uint32_t mapctl_buf_length;
753     uint32_t mapctl_buf_data_len;
754     int64_t mapctl_agg_offset;
755     int8_t mapctl_buf_write;
756     char aggrec_rsvdk[3]; /* pad to quadword boundary */
757
758     /* ***** MAP LEAF I/O BUFFER ***** */
759
760     #define MAPLEAF_IO_BUFSIZE \
761         (((XTPAGE_SIZE*4)+BYTESPERPAGE-1)/BYTESPERPAGE)*BYTESPERPAGE
762
763     char maplf_info_eyecatcher[8];
764     char *mapleaf_buf_ptr;
765     uint32_t mapleaf_buf_length;
766     uint32_t mapleaf_buf_data_len;
767     int64_t mapleaf_agg_offset;
768     int8_t mapleaf_for_aggregate;
769     int8_t mapleaf_which_it;
770     int8_t mapleaf_buf_write;
771     char aggrec_rsvdm[1]; /* pad to quadword boundary */
772
773     /* ***** BLOCK MAP LEVEL PAGE I/O BUFFER ***** */
774
775     #define BMAPLV_IO_BUFSIZE BYTESPERPAGE
776
777     char bmaplv_info_eyecatcher[8];
778     char *bmaplv_buf_ptr;
779     uint32_t bmaplv_buf_length;
780     uint32_t bmaplv_buf_data_len;
781     int64_t bmaplv_agg_offset;
782     int64_t bmaplv_logical_offset;
783     int8_t bmaplv_current_level;
784     int8_t bmaplv_buf_write;
785     char aggrec_rsvdo[10]; /* pad to quadword boundary */
786
787     /* ***** BLOCK MAP DMAP PAGE I/O BUFFER ***** */
788     /*
789     * note: This is not actually a unique buffer. Since
790     * the Inode Allocation Maps verification is
791     * completed before the Block Allocation Map
792     * verification, buffer space is allocated and
793     * used first for Inode Alloc Map verification
794     * processing and then for Block Allocation Map
795     * verification.
796     */
797
798     #define BMAPDMP_IO_BUFSIZE IAG_IO_BUFSIZE
799
800     char bmpdm_info_eyecatcher[8];
801     char *bmpdm_buf_ptr;
802     uint32_t bmpdm_buf_length;
803     uint32_t bmpdm_buf_data_len;
804     int64_t bmpdm_agg_offset;
805     int64_t bmpdm_logical_offset;
806     int8_t bmpdm_buf_write;
807     char aggrec_rsvdq[11]; /* pad to quadword boundary */
808
809     /* ***** INODE I/O BUFFER ***** */
810

```

```

811 #define INODE_IO_BUFSIZE \
812     (((INODE_EXTENT_SIZE+BYTESPERPAGE-1)/BYTESPERPAGE)*BYTESPERPAGE)
813
814     char inobuf_info_eyecatcher[8];
815     char *ino_buf_ptr;
816     uint32_t ino_buf_length;
817     uint32_t ino_buf_data_len;
818     int64_t ino_buf_agg_offset;      /* agg byte offset of buf contents */
819     uint32_t ino_buf_1st_ino;      /* 1st inode number in buffer */
820     uint32_t ino_fsnum;
821     pxd_t ino_ixpxd;
822     int8_t ino_for_aggregate;
823     int8_t ino_which_it;
824     int8_t ino_buf_write;
825     char aggrec_rsvds[1]; /* pad to quadword boundary */
826
827     /* ***** INTERNAL/LEAF NODE I/O BUFFER ***** */
828
829 #define NODE_IO_BUFSIZE \
830     (((XTPAGE_SIZE*4)+BYTESPERPAGE-1)/BYTESPERPAGE)*BYTESPERPAGE)
831
832     char nodbuf_info_eyecatcher[8];
833     char *node_buf_ptr;
834     uint32_t node_buf_length;
835     uint32_t node_buf_data_len;
836     int64_t node_agg_offset;
837     int8_t node_buf_write;
838     char aggrec_rsvdu[3]; /* pad to quadword boundary */
839
840     /* ***** instantiation of imap AG tables ***** */
841
842     char agg_AGTbl_eyecatcher[8];
843     char aggrec_rsvdw[8]; /* pad to quadword boundary */
844     struct fsck_ag_record agg_AGTbl[MAXAG];
845
846     char fset_AGTbl_eyecatcher[8];
847     char aggrec_rsvdy[8]; /* pad to quadword boundary */
848     struct fsck_ag_record fset_AGTbl[MAXAG];
849
850     /* ***** instantiation of imap IAG workspace ***** */
851     char amap_eyecatcher[8];
852     uint32_t amap[EXTSPERIAG]; /* 512 : inode allocation map */
853     char aggrec_rsvd00[8]; /* pad to quadword boundary */
854     char fextsumm_eyecatcher[8];
855     int32_t fextsumm[SMAPSZ]; /* 16 : free extent summary map */
856     char aggrec_rsvd02[8]; /* pad to quadword boundary */
857     char finosumm_eyecatcher[8];
858     int32_t finosumm[SMAPSZ]; /* 16 : free inode summary map */
859
860     /* ***** instantiation of blockmap workspace ***** */
861
862     struct blkmap_wspace blkmap_wsp;
863 };
864
865 #endif

```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_INODE
19 #define H_INODE
20
21 #include <jfs_types.h>
22 #include "devices.h"
23
24 int ujfs_rwinode(HFILE, struct dinode *, uint32_t, int32_t, int32_t, uint32_t, uint32_t);
25 int ujfs_rwddr(HFILE, int64_t *, struct dinode *, int64_t, int32_t, int32_t);
26
27 #endif
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_UJFS_ENDIAN
19 #define _H_UJFS_ENDIAN
20
21 #include "jfs_types.h"
22 #include "jfs_byteorder.h"
23 #include "jfs_superblock.h"
24 #include "jfs_dmap.h"
25 #include "jfs_imap.h"
26 #include "jfs_dinode.h"
27 #include "jfs_logmgr.h"
28 #include "fsckwsp.h"
29
30 #if __BYTE_ORDER == __BIG_ENDIAN
31 void ujfs_swap_dbmap(struct dbmap *);
32 void ujfs_swap_dinode(struct dinode *, int32_t, uint32_t);
33 void ujfs_swap_dinomap(struct dinomap *);
34 void ujfs_swap_dmap(struct dmap *);
35 void ujfs_swap_dmapctl(struct dmapctl *);
36 void ujfs_swap_dtpage_t(dtpage_t *, uint32_t);
37 void ujfs_swap_fsck_blk_map_hdr(struct fsck_blk_map_hdr *);
38 void ujfs_swap_fsck_blk_map_page(struct fsck_blk_map_page *);
39 void ujfs_swap_fscklog_entry_hdr(struct fscklog_entry_hdr *);
40 void ujfs_swap_iag(struct iag *);
41 void ujfs_swap_logpage(struct logpage *, uint8_t);
42 void ujfs_swap_logsuper(struct logsuper *);
43 void ujfs_swap_lrd(struct lrd *);
44 void ujfs_swap_superblock(struct superblock *);
45 void ujfs_swap_xtpage_t(xtpage_t *);
46
47 static inline void ujfs_swap_inoext(struct dinode *ptr, int32_t mode, uint32_t flag)
48 {
49     int i;
50     for (i = 0; i < INOSPEREXT; i++, ptr++)
51         ujfs_swap_dinode(ptr, mode, flag);
52 }
53
54 #define swap_multiple(swap_func, ptr, num) \
55 do { \
56     int i; \
57     for (i = 0; i < num; i++, (ptr)++) \
58         swap_func(ptr); \
59 } while (0)
60
61 #else /* Little endian */
62
63 #define ujfs_swap_dbmap(dbmap) do { while (0)
64 #define ujfs_swap_dinode(dinode, mode, flag) do { while (0)
65 #define ujfs_swap_dinomap(dinomap) do { while (0)
66 #define ujfs_swap_dmap(dmap) do { while (0)
67 #define ujfs_swap_dmapctl(dmapctl) do { while (0)
68 #define ujfs_swap_dtpage_t(dtpage, flag) do { while (0)
69 #define ujfs_swap_fsck_blk_map_hdr(map_hdr) do { while (0)
70 #define ujfs_swap_fsck_blk_map_page(map_page) do { while (0)
71 #define ujfs_swap_fscklog_entry_hdr(ent_hdr) do { while (0)
72 #define ujfs_swap_iag(iag) do { while (0)
73 #define ujfs_swap_inoext(inoext, mode, flag) do { while (0)
74 #define ujfs_swap_logpage(logpage, pages) do { while (0)
75 #define ujfs_swap_logsuper(logsuper) do { while (0)
76 #define ujfs_swap_lrd(lrd) do { while (0)
77 #define ujfs_swap_superblock(sb) do { while (0)
78 #define ujfs_swap_xtpage_t(xtpage) do { while (0)
79 #define swap_multiple(swap_func, ptr, num) do { while (0)
80
81 #endif
82
83 #endif /* _H_UJFS_ENDIAN */

```

```
1 /*
2  * Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  * This program is free software; you can redistribute it and/or modify
5  * it under the terms of the GNU General Public License as published by
6  * the Free Software Foundation; either version 2 of the License, or
7  * (at your option) any later version.
8  *
9  * This program is distributed in the hope that it will be useful,
10 * but WITHOUT ANY WARRANTY; without even the implied warranty of
11 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 * the GNU General Public License for more details.
13 *
14 * You should have received a copy of the GNU General Public License
15 * along with this program; if not, write to the Free Software
16 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_LIBFS
19 #define H_LIBFS
20
21 #define LIBFS_BADMAGIC -5 /* magic number not recognized */
22 #define LIBFS_BADVERSION -6 /* magic num ok, incompatible vers */
23 #define LIBFS_CORRUPTSUPER -10 /* fragsize, agsize, or iagsize bad */
24
25 #endif /* H_LIBFS */
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_LOGFORM
19 #define H_LOGFORM
20
21 int jfs_logform(int, int, int, uint, int64_t, int, uid_t, char *);
22
23 #endif          /* H_LORFORM */
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #include <stdio.h>
19 #include <stdlib.h>
20 #include <string.h>
21
22 #include "jfs_types.h"
23 #include "utilsubs.h"
24
25 /*
26 *   log2shift()
27 */
28 int32_t log2shift(uint32_t n)
29 {
30     uint32_t shift = 0;
31
32     while (n > 1) {
33         /* n is not power of 2 */
34         if (n & 1)
35             return -1;
36
37         shift++;
38         n >>= 1;
39     }
40
41     return shift;
42 }
43
44 /*
45 *   ui
46 *   ==
47 */
48 /*
49 *   prompt()
50 */
51 char prompt(char *str)
52 {
53     char cmd[81];
54
55     fputs(str, stdout);
56     fflush(stdout);
57
58     /* get NULL terminated input */
59     fgets(cmd, 81, stdin);
60
61     return cmd[0];          /* return response letter */
62 }
63
64 /*
65 *   more()
66 */
67 int more(void)
68 {
69     char cmd[81];
70
71     fputs("- hit Enter to continue, e[x]it -", stdout);
72     fflush(stdout);
73
74     /* get NULL terminated input */
75     fgets(cmd, 80, stdin);
76
77     if (cmd[0] == 'x')
78         return 1;          /* do NOT continue */
79     else
80         return 0;          /* continue */
81 }
```

```
1 /*
2  * Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  * This program is free software; you can redistribute it and/or modify
5  * it under the terms of the GNU General Public License as published by
6  * the Free Software Foundation; either version 2 of the License, or
7  * (at your option) any later version.
8  *
9  * This program is distributed in the hope that it will be useful,
10 * but WITHOUT ANY WARRANTY; without even the implied warranty of
11 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 * the GNU General Public License for more details.
13 *
14 * You should have received a copy of the GNU General Public License
15 * along with this program; if not, write to the Free Software
16 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef _H_UJFS_UTILSUBS
19 #define _H_UJFS_UTILSUBS
20
21 /*
22  * utilsubs.h
23  */
24
25 /*
26  * function prototypes
27  */
28 int32_t log2shift(uint32_t n);
29 char prompt(char *str);
30 int more(void);
31
32 #endif /* _H_UJFS_UTILSUBS */
```



```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 /*
19 *   FUNCTIONS: no-frills substitutes for fsck routines
20 *               used by logredo modules outside of fsck
21 */
22 #include <stdio.h>
23 #include <stdlib.h>
24 #include <string.h>
25
26 #include "jfs_types.h"
27 #include "fsckmsgc.h"
28 #include "fsckmsgc.h"
29
30 #define STDOUT_HANDLE 1
31 #define STDERR_HANDLE 2
32
33 extern short MsgProtocol[][2];
34
35 extern char *Vol_Label;
36
37 extern char *verbose_msg_ptr;
38
39 extern char *msgprms[];
40 extern short msgprmidx[];
41
42 extern unsigned long msgs_txt_maxlen;
43
44 /*****
45 * NAME: alloc_wrksp
46 *
47 * FUNCTION: Allocates and initializes (to guarantee the storage is backed)
48 *           dynamic storage for the caller.
49 *
50 * PARAMETERS:
51 *   length          - input - the number of bytes of storage which are needed
52 *   dynstg_object   - input - a constant (see xfscck.h) identifying the purpose
53 *                   for which the storage is needed (Used in error
54 *                   message if the request cannot be satisfied.
55 *   addr_wrksp_ptr  - input - the address of a variable in which this routine
56 *                   will return the address of the dynamic storage
57 *                   allocated for the caller
58 *
59 * NOTES:
60 *
61 * RETURNS:
62 *   success: 0
63 *   failure: something else
64 */
65 int alloc_wrksp(unsigned length, int dynstg_object, int for_logredo, void **addr_wrksp_ptr)
66 {
67     int awsp_rc = 0;
68     unsigned min_length;
69
70     *addr_wrksp_ptr = NULL; /* initialize return value */
71     min_length = ((length + 7) / 4) * 4; /* round up to an 4 byte boundary */
72
73     *addr_wrksp_ptr = (char *) malloc(min_length);
74
75     return (awsp_rc);
76 } /* end alloc_wrksp() */
77
78 /*****
79 * NAME: fsck_send_msg
80 *
81 * FUNCTION: Issue an fsck message, depending on the message's protocol
82 *           according to the fsck message arrays.
83 *
84 * PARAMETERS:
85 *   ?               - input -
86 *   ?               - returned -
87 *
88 * NOTES: none
89 *
90 * RETURNS:

```

```
91  * nothing
92  */
93  void fsck_send_msg(int msg_num, int num_inserts)
94  {
95      int prmidx;
96
97      memset((void *) verbose_msg_ptr, '\0', msgs_txt_maxlen);
98
99      prmidx = 0;
100     while (prmidx < num_inserts) {
101         if (msgprmidx[prmidx] != 0) { /* variable text insert */
102             msgprms[prmidx] = MsgText[msgprmidx[prmidx]];
103         } /* end variable text insert */
104         prmidx += 1;
105     }
106
107     sprintf(verbose_msg_ptr, MsgText[msg_num],
108            msgprms[0], msgprms[1], msgprms[2], msgprms[3], msgprms[4],
109            msgprms[5], msgprms[6], msgprms[7], msgprms[8], msgprms[9]);
110
111     printf("%s", verbose_msg_ptr);
112
113     return;
114 } /* end fsck_send_msg() */
```

```
1 /*
2  * Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  * This program is free software; you can redistribute it and/or modify
5  * it under the terms of the GNU General Public License as published by
6  * the Free Software Foundation; either version 2 of the License, or
7  * (at your option) any later version.
8  *
9  * This program is distributed in the hope that it will be useful,
10 * but WITHOUT ANY WARRANTY; without even the implied warranty of
11 * MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 * the GNU General Public License for more details.
13 *
14 * You should have received a copy of the GNU General Public License
15 * along with this program; if not, write to the Free Software
16 * Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_INITMAP
19 #define H_INITMAP
20
21 #define ALLOC          0x1
22 #define FREE          0x2
23 #define BADBLOCK      0x4
24
25 int calc_map_size(int64_t, struct dinode *, int, int *, unsigned);
26 int markit(int, unsigned);
27 int record_LVM_BadBlks( int, int, int, struct dinode *, int64_t );
28 int verify_last_blocks( int, int, struct dinode * );
29 int write_block_map(int, int64_t, int);
30
31 #endif /* H_INITMAP */
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #ifndef H_INODEMAP
19 #define H_INODEMAP
20
21 #include "devices.h"
22
23 int init_inode_map(int, HFILE, int64_t, int, int64_t, int, unsigned short, int, unsigned);
24
25 #endif
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18
19 /* JFS includes */
20 #include "jfs_filsys.h"
21 #include "jfs_superblock.h"
22 #include "jfs_logmgr.h"
23
24 /*-----
25 * NAME: build_flag_string
26 *
27 * FUNCTION: build a string of text descriptions of
28 *           the set flags in the superblock
29 *
30 * PARAMETERS:
31 *   flags           - flags field from the superblock
32 *   flag_string     - returned string of flag descriptions
33 */
34 void build_flag_string(uint32_t flags, char *flag_string)
35 {
36     char *string_ptr = flag_string;
37
38     if (flags & JFS_LINUX) {
39         strcpy(string_ptr, "JFS_LINUX ");
40         string_ptr += 11;
41     }
42     if (flags & JFS_OS2) {
43         strcpy(string_ptr, "JFS_OS2 ");
44         string_ptr += 9;
45     }
46     if (flags & JFS_COMMIT) {
47         strcpy(string_ptr, "JFS_COMMIT ");
48         string_ptr += 12;
49     }
50     if (flags & JFS_GROUPCOMMIT) {
51         strcpy(string_ptr, "JFS_GROUPCOMMIT ");
52         string_ptr += 17;
53     }
54     if (flags & JFS_LAZYCOMMIT) {
55         strcpy(string_ptr, "JFS_LAZYCOMMIT ");
56         string_ptr += 16;
57     }
58     if (flags & JFS_INLINELOG) {
59         strcpy(string_ptr, "JFS_INLINELOG ");
60         string_ptr += 15;
61     }
62     if (flags & JFS_BAD_SAIT) {
63         strcpy(string_ptr, "JFS_BAD_SAIT ");
64         string_ptr += 14;
65     }
66     if (flags & JFS_SPARSE) {
67         strcpy(string_ptr, "JFS_SPARSE ");
68         string_ptr += 12;
69     }
70     if (flags & JFS_DASD_ENABLED) {
71         strcpy(string_ptr, "JFS_DASD_ENABLED ");
72         string_ptr += 18;
73     }
74     if (flags & JFS_DASD_PRIME) {
75         strcpy(string_ptr, "JFS_DASD_PRIME ");
76         string_ptr += 16;
77     }
78     strcpy(string_ptr, "\0");
79
80     return;
81 }
82
83 /*-----
84 * NAME: display_super
85 *
86 * FUNCTION: display the JFS file system superblock
87 *
88 * PARAMETERS:
89 *   sb - pointer to the file system superblock
90 *

```

```

91  * SAMPLE OUTPUT:
92
93  JFS magic number:      'JFS1'
94  JFS version:          2
95  JFS state:            clean
96  JFS flags:            JFS_LINUX JFS_COMMIT JFS_GROUPCOMMIT
97  Aggregate block size: 4096 bytes
98  Aggregate size:       2055784 blocks
99  Physical block size:  512 bytes
100 Allocation group size: 8192 aggregate blocks
101 Log device number:    0x3f06
102 Filesystem creation:  Fri Aug  2 14:48:53 2002
103 File system UUID:     7940a3b8-0bb0-4c0d-a389-630619783c6e
104 Volume label:        ' '
105 External log UUID:    47c80a1d-7f25-4c61-9596-f33e463e5b38
106
107  */
108 void display_super(struct superblock *sb)
109 {
110     char *state;
111     char uuid_unparsed[37];
112     char flag_string[142];
113     time_t tm;
114
115     switch (sb->s_state) {
116     case FM_CLEAN:
117         state = "clean";
118         break;
119     case FM_MOUNT:
120         state = "mounted";
121         break;
122     case FM_DIRTY:
123         state = "dirty";
124         break;
125     case FM_LOGREDO:
126         state = "logredo";
127         break;
128     default:
129         state = "unknown";
130         break;
131     }
132
133     build_flag_string(sb->s_flag, flag_string);
134
135     printf("\nJFS filesystem superblock:\n\n");
136     printf("JFS magic number:\t'%4.4s'\n", sb->s_magic);
137     printf("JFS version:\t\t%d\n", sb->s_version);
138     printf("JFS state:\t\t%s\n", state);
139     printf("JFS flags:\t\t%s\n", flag_string);
140     printf("Aggregate block size:\t%d bytes\n", sb->s_bsize);
141     printf("Aggregate size:\t\t%lld blocks\n", (long long) sb->s_size);
142     printf("Physical block size:\t%d bytes\n", sb->s_pbsize);
143     printf("Allocation group size:\t%d aggregate blocks\n", sb->s_agsize);
144     printf("Log device number:\t0x%x\n", sb->s_logdev);
145     tm = sb->s_time.tv_sec;
146     printf("Filesystem creation:\t%s", ctime(&tm));
147     /*
148     * JFS version 2 incorporated new fields for
149     * UUID, log UUID, and a 16 char volume label
150     * into the superblock. Account for those here.
151     */
152     if (sb->s_version == JFS_VERSION) {
153         uuid_unparse(sb->s_uuid, uuid_unparsed);
154         printf("File system UUID:\t%s\n", uuid_unparsed);
155         printf("Volume label:\t\t'%16s'\n", sb->s_label);
156         uuid_unparse(sb->s_loguuid, uuid_unparsed);
157         printf("External log UUID:\t%s\n", uuid_unparsed);
158     } else {
159         printf("Volume label:\t\t'%11s'\n", sb->s_fpack);
160     }
161     printf("\n");
162
163     return;
164 }
165
166 /*-----
167  * NAME: display_logsuper
168  * FUNCTION: display the JFS log superblock
169  *
170  * PARAMETERS:
171  *     lsp - pointer to the log superblock
172  *
173  * SAMPLE OUTPUT:
174
175  JFS log magic number:  0x87654321
176  JFS log version:      1
177  Log opened/mounted:   3
178  Logical block size:   4096 bytes
179  Log size:             32768 blocks

```

```

181 Log flags:                JFS_LINUX JFS_COMMIT JFS_GROUPCOMMIT
182 Log state:                LOGMOUNT
183 Log device UUID:         47c80a1d-7f25-4c61-9596-f33e463e5b38
184 Log volume label:       ' JFSLogVol'
185 Active file systems:
186     active[0]: 7940a3b8-0bb0-4c0d-a389-630619783c6e
187
188 */
189 void display_logsuper(struct logsuper * lsp)
190 {
191     char *state;
192     char uuid_unparsed[37];
193     int i;
194     bool active_fs = false;
195     char flag_string[142];
196
197     switch (lsp->state) {
198     case LOGMOUNT:
199         state = "LOGMOUNT";
200         break;
201     case LOGREDONE:
202         state = "LOGREDONE";
203         break;
204     case LOGWRAP:
205         state = "LOGWRAP";
206         break;
207     case LOGREADERR:
208         state = "LOGREADERR";
209         break;
210     default:
211         state = "Unknown";
212         break;
213     }
214
215     build_flag_string(lsp->flag, flag_string);
216
217     printf("\nJFS external log superblock:\n\n");
218     printf("JFS log magic number:\t0x%x\n", lsp->magic);
219     printf("JFS log version:\t%d\n", lsp->version);
220     printf("Log opened/mounted:\t%d\n", lsp->serial);
221     printf("Logical block size:\t%d bytes\n", lsp->bsize);
222     printf("Log size:\t%d blocks\n", lsp->size);
223     printf("Log flags:\t%s\n", flag_string);
224     printf("Log state:\t%s\n", state);
225     uuid_unparse(lsp->uuid, uuid_unparsed);
226     printf("Log device UUID:\t%s\n", uuid_unparsed);
227     printf("Log volume label:\t'%16s'\n", lsp->label);
228     printf("Active file systems:");
229     for (i = 0; i < MAX_ACTIVE; i++) {
230         uuid_unparse(lsp->active[i], uuid_unparsed);
231         /* only print active (non-zero uuid) file systems */
232         if (strcmp(uuid_unparsed, "00000000-0000-0000-0000-000000000000", 36)) {
233             printf("\n active[%d]: %s", i, uuid_unparsed);
234             active_fs = true;
235         }
236     }
237
238     if (!active_fs) {
239         printf("\tNone active.\n");
240     }
241     printf("\n");
242
243     return;
244 }

```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #include "config.h"
19
20 #include <fcntl.h>
21 #include <unistd.h>
22 #include <sys/stat.h>
23 #include <string.h>
24 #include <stdlib.h>
25 #include <stdio.h>
26 #include <time.h>
27 #include <ctype.h>
28 #include "jfs_types.h"
29 #include "jfs_endian.h"
30 #include "jfs_filsys.h"
31 #include "jfs_superblock.h"
32 #include "inode.h"
33 #include "super.h"
34 #include "jfs_version.h"
35
36 #define EXIT(fd, rc) {close(fd); exit(rc);}
37
38 static int J_flag, l_flag, L_flag, U_flag;
39 char *new_label, *new_UUID;
40 char *device;
41 char logdev[255] = { '\0' };
42 int log_fd = -1;
43
44 extern int LogOpenMode;
45 extern int open_by_label(uuid_t, int, int, char *);
46
47 void tune_usage()
48 {
49     printf( "\nUsage: jfs_tune [-J options] [-l] [-L vol_label] [-U uuid] [-V] device\n"
50           "\nEmergency help:\n"
51           "  -J options  Set external journal options.\n"
52           "  -l         Display superblock\n"
53           "  -L vol_label Set volume label.\n"
54           "  -U uuid    Set UUID.\n"
55           "  -V         Print version information only.\n" );
56     exit(-1);
57 }
58
59 /*-----
60 * NAME: parse_journal_opts
61 *
62 * FUNCTION: parse journal (-J) options
63 *           set log file descriptor (global log_fd)
64 *           set log device name (global logdev)
65 *
66 * PARAMETERS:
67 *   opts - options string
68 */
69 void parse_journal_opts(const char *opts)
70 {
71     int journal_usage = 0;
72     uuid_t log_uuid;
73
74     LogOpenMode = O_RDWR | O_EXCL;
75
76     if (strncmp(opts, "device=", 7) == 0) {
77         if (strncmp(opts + 7, "UUID=", 5) == 0) {
78             if (uuid_parse((char *) opts + 7 + 5, log_uuid)) {
79                 fputs("\nError: UUID entered in improper format.\n",
80                       stderr);
81                 exit(-1);
82             } else {
83                 log_fd = open_by_label(log_uuid, 0, 1, logdev);
84             }
85         } else if (strncmp(opts + 7, "LABEL=", 6) == 0) {
86             log_fd = open_by_label((char *) opts + 7 + 6, 1, 1, logdev);
87         } else {
88             strcpy(logdev, ((char *) opts + 7));
89             if (logdev) {
90                 log_fd = open(logdev, LogOpenMode, 0);

```



```

91         } else {
92             journal_usage++;
93         }
94     }
95     } else
96         journal_usage++;
97
98     if (journal_usage) {
99         fprintf(stderr, "\nInvalid journal option(s) specified.\n\n"
100             "Valid options for -J are:\n"
101             "\tdevice=<journal device>\n"
102             "\tdevice=UUID=<UUID of journal device>\n"
103             "\tdevice=LABEL=<label of journal device>\n\n");
104         exit(1);
105     }
106
107     return;
108 }
109
110 /*-----
111  * NAME: parse_tune_options
112  *
113  * FUNCTION: parse tune options
114  *
115  * PARAMETERS:
116  *     argc - number of passed arguments
117  *     argv - string of arguments
118  *
119  * RETURNS:
120  *     success: 0
121  *     failure: any other value
122  */
123 static void parse_tune_options(int argc, char *argv[])
124 {
125     int c;
126
127     while ((c = getopt(argc, argv, "J:IL:U:V")) != EOF) {
128         switch (c) {
129             case 'J':
130                 /* attach external journal device */
131                 parse_journal_opts(optarg);
132                 J_flag = 1;
133                 break;
134             case 'l':
135                 /* display superblock */
136                 l_flag = 1;
137                 break;
138             case 'L':
139                 /* set volume label */
140                 new_label = optarg;
141                 L_flag = 1;
142                 break;
143             case 'U':
144                 /* set UUID */
145                 new_UUID = optarg;
146                 U_flag = 1;
147                 break;
148             case 'V':
149                 /* print version and exit */
150                 exit(0);
151                 break;
152             default:
153                 tune_usage();
154                 break;
155         }
156     }
157
158     if (optind != argc - 1) {
159         printf("\nError: Device not specified or command format error.\n");
160         tune_usage();
161     }
162
163     if (!J_flag && !l_flag && !L_flag && !U_flag) {
164         printf("\nError: No options selected.\n");
165         tune_usage();
166     }
167
168     device = argv[optind];
169
170     return;
171 }
172
173 /*-----
174  * NAME: main
175  *
176  * FUNCTION: adjust JFS tunable parameters
177  *
178  * PARAMETERS:
179  *     argc - number of passed arguments
180  *     argv - string of arguments

```

```

181  *
182  * RETURNS:
183  *   success: 0
184  *   failure: any other value
185  */
186  int main(int argc, char *argv[])
187  {
188      char *argp;
189      int rc = 0;
190      int fd, superblock_type;
191      bool mounted = false;
192      struct superblock sb;
193      struct logsuper logsup;
194
195      #define FS_SUPER_SECONDARY 0
196      #define FS_SUPER_PRIMARY 1
197      #define LOG_SUPER 2
198
199      printf("jfs_tune version %s,%s\n", VERSION, JFSUTILS_DATE);
200
201      parse_tune_options(argc, argv);
202
203      /*
204       * Check if device is mounted. -l is the only parameter
205       * supported on a mounted device, so if any others are
206       * selected, let the user know. If the device is mounted
207       * and -l was not specified, get out.
208       */
209      if (rc = Is_Device_Mounted(device)) {
210          mounted = true;
211          if (J_flag || L_flag || U_flag) {
212              fprintf(stderr, "\n%s is mounted.\n"
213                  "While mounted, the only supported jfs_tune parameter is -l.\n",
214                  device);
215              if (!l_flag) {
216                  exit(-1);
217              }
218          }
219      }
220
221      /* Open device */
222      fd = open(device, O_RDWR | O_EXCL, 0);
223      if (fd < 0) {
224          fprintf(stderr, "Error: Cannot open device %s.\n", device);
225          exit(-1);
226      }
227
228      /* Get and validate primary JFS superblock */
229      if ((rc = ujfs_get_superblk(fd, &sb, 1)) == 0) {
230          if ((rc = ujfs_validate_super(&sb)) == 0) {
231              superblock_type = FS_SUPER_PRIMARY;
232          }
233      }
234
235      /* If failure retrieving primary superblock, get/validate secondary superblock */
236      if (rc) {
237          if ((rc = ujfs_get_superblk(fd, &sb, 0)) == 0) {
238              if ((rc = ujfs_validate_super(&sb)) == 0) {
239                  superblock_type = FS_SUPER_SECONDARY;
240              }
241          }
242      }
243
244      /* If no valid FS superblock, see if we have a log superblock */
245      if (rc) {
246          if ((rc = ujfs_get_logsuper(fd, &logsup)) == 0) {
247              if ((rc = ujfs_validate_logsuper(&logsup)) == 0) {
248                  superblock_type = LOG_SUPER;
249                  /*
250                   * We know this is an external journal device.
251                   * Now check to see if it is attached to a
252                   * mounted file system. If so, the only
253                   * valid option is -l.
254                   */
255                  if (logsup.state == LOGMOUNT) {
256                      mounted = true;
257                      if (J_flag || L_flag || U_flag) {
258                          fprintf(stderr,
259                              "\n%s contains an external journal for a mounted filesystem.\n"
260                              "While mounted, the only supported jfs_tune parameter is -l.\n",
261                              device);
262                          if (!l_flag) {
263                              EXIT(fd, -1);
264                          }
265                      }
266                  }
267              }
268          }
269      }
270

```

```

271     /* If we couldn't find/read a valid JFS FS or log superblock, warn user and exit */
272     if (rc) {
273         printf("\nCould not read valid JFS FS or log superblock on device %s.\n",
274             device);
275         EXIT(fd, -1);
276     }
277
278     /*
279     * Account for bug in mkfs.jfs 1.0.18 and 1.0.19
280     * that didn't properly set the file system version
281     * number to 2 when using an external journal.
282     */
283     if (!mounted && (superblock_type < LOG_SUPER) &&
284         (sb.s_version < JFS_VERSION) && !(sb.s_flag & JFS_INLINELOG)) {
285         sb.s_version = JFS_VERSION;
286         rc = ujfs_put_superblk(fd, &sb, superblock_type);
287         if (rc) {
288             printf("\nCould not update JFS version number properly on %s.\n",
289                 device);
290             EXIT(fd, rc);
291         }
292     }
293
294     /*
295     * set volume label on unmounted device
296     */
297     if (L_flag && !mounted) {
298         if (superblock_type < LOG_SUPER) {
299             /* change label in JFS file system superblock */
300             /*
301             * The superblock in JFS releases before 1.0.18 stores
302             * the label in s_fpack[11]. The superblock in JFS
303             * releases 1.0.18 and greater has s_fpack, but uses
304             * the new field s_label[16] to store the label.
305             * s_label is in an area of the superblock that was
306             * allocated but unused in pre 1.0.18, so if per chance
307             * the user is using an old JFS file system, setting
308             * s_label will not be a problem.
309             */
310             memset(sb.s_fpack, 0, sizeof (sb.s_fpack));
311             strncpy(sb.s_fpack, new_label, sizeof (sb.s_fpack));
312             if (strlen(new_label) > sizeof (sb.s_label))
313                 fprintf(stderr, "Warning: label too long, truncating.\n");
314             memset(sb.s_label, 0, sizeof (sb.s_label));
315             strncpy(sb.s_label, new_label, sizeof (sb.s_label));
316             rc = ujfs_put_superblk(fd, &sb, superblock_type);
317         } else {
318             /* change label in JFS log superblock */
319             if (strlen(new_label) > sizeof (logsup.label))
320                 fprintf(stderr, "Warning: label too long, truncating.\n");
321             memset(logsup.label, 0, sizeof (logsup.label));
322             strncpy(logsup.label, new_label, sizeof (logsup.label));
323             rc = ujfs_put_logsuper(fd, &logsup);
324         }
325         if (rc) {
326             printf("\nError writing superblock to disk. Label unchanged.\n");
327             EXIT(fd, rc);
328         } else {
329             printf("Volume label updated successfully.\n");
330         }
331     }
332
333     /*
334     * set UUID on unmounted device
335     */
336     if (U_flag && !mounted) {
337         uuid_t *uu;
338         uu = ((superblock_type < LOG_SUPER) ? &sb.s_uuid : &logsup.uuid);
339         if ((strcasecmp(new_UUID, "null") == 0) ||
340             (strcasecmp(new_UUID, "clear") == 0)) {
341             uuid_clear(*uu);
342         } else if (strcasecmp(new_UUID, "time") == 0) {
343             uuid_generate_time(*uu);
344         } else if (strcasecmp(new_UUID, "random") == 0) {
345             uuid_generate(*uu);
346         } else if (uuid_parse(new_UUID, *uu)) {
347             fprintf(stderr, "Invalid UUID format.\n");
348             EXIT(fd, -1);
349         }
350
351         if (superblock_type < LOG_SUPER) {
352             rc = ujfs_put_superblk(fd, &sb, superblock_type);
353         } else {
354             rc = ujfs_put_logsuper(fd, &logsup);
355         }
356
357         if (rc) {
358             printf("\nError writing superblock to disk. UUID unchanged.\n");
359             EXIT(fd, rc);
360         } else {

```

```

361         printf("UUID updated successfully.\n");
362     }
363 }
364
365 /*
366  * attach external journal to JFS file system
367  */
368 if (J_flag && !mounted) {
369     /*
370      * NOTE: If we ever allow attaching more than one file system to a
371      * single log file, we'll have to change the conditions of the above
372      * 'if' to account for a log file that is in use by one file system
373      * (state LOGMOUNT), but is being attached to by another file system.
374      */
375     struct stat st;
376
377     /* make sure device to be attached to is a JFS file system */
378     if (superblock_type >= LOG_SUPER) {
379         printf("\nError: %s does not contain a JFS file system.\n", device);
380         EXIT(fd, -1);
381     }
382
383     /* log_fd was set in parse_journal_opts */
384     if (log_fd < 0) {
385         printf("\nError: Could not find/open specified external journal device.\n");
386         EXIT(fd, -1);
387     }
388
389     /* get valid log superblock */
390     if ((rc = ujfs_get_logsuper(log_fd, &logsup)) == 0) {
391         if ((rc = ujfs_validate_logsuper(&logsup)) == 0) {
392             if (fstat(log_fd, &st) {
393                 rc = -1;
394             } else {
395                 /* update FS superblock */
396                 sb.s_logdev = dev_to_kdev(st.st_rdev);
397                 uuid_copy(sb.s_loguuid, logsup.uuid);
398                 sb.s_version = JFS_VERSION;
399                 sb.s_flag &= (~JFS_INLINELOG);
400                 memset(&sb.s_logpxd, 0, sizeof (pxd_t));
401                 rc = ujfs_put_superblk(fd, &sb, superblock_type);
402             }
403         }
404     }
405
406     /* If we could't find/read a valid JFS log superblock, let user know */
407     if (rc) {
408         printf("\nError attaching JFS external journal to JFS FS.\n");
409     } else {
410         printf("Attached JFS external journal to JFS FS successfully.\n");
411     }
412     close(log_fd);
413 }
414
415 /*
416  * display superblock
417  */
418 if (l_flag) {
419     if (superblock_type < LOG_SUPER) {
420         display_super(&sb);
421     } else {
422         display_logsuper(&logsup);
423     }
424 }
425
426 close(fd);
427 return rc;
428 }

```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 #include "xpeek.h"
19
20 /* libfs includes */
21 #include "devices.h"
22 #include "utilsubs.h"
23
24 void alter()
25 {
26     int64_t address;
27     int64_t block;
28     uint8_t *buffer;
29     uint8_t byte;
30     char cmdline[512];
31     char *hexstring;
32     unsigned hex_length;
33     unsigned length;
34     unsigned offset;
35     char *ptr;
36     char *token;
37
38     token = strtok(0, " \n");
39     if (token == 0) {
40         fputs("alter: Please enter: block offset hex-digits> ", stdout);
41         fgets(cmdline, 512, stdin);
42         token = strtok(cmdline, " \n");
43         if (token == 0)
44             return;
45     }
46     errno = 0;
47     block = strtoull(token, 0, 0);
48     if (block == 0 && errno) {
49         fputs("alter: invalid block\n\n", stderr);
50         return;
51     }
52     address = block << l2bsize;
53
54     token = strtok(0, " \n");
55     if (!token) {
56         fputs("alter: Not enough arguments!\n", stderr);
57         return;
58     }
59     offset = strtoul(token, 0, 16);
60     if (offset == 0 && errno) {
61         fputs("alter: invalid offset\n", stderr);
62         return;
63     }
64     hexstring = strtok(0, " \n");
65     if (!hexstring) {
66         fputs("alter: Not enough arguments!\n", stderr);
67         return;
68     }
69     if (strtok(0, " \n")) {
70         fputs("alter: Too many arguments!\n", stderr);
71         return;
72     }
73     hex_length = strlen(hexstring);
74     if (hex_length & 1) {
75         /* odd number of hex digits */
76         fputs("alter: hex string must have even number of digits!\n", stderr);
77         return;
78     }
79
80     /* Round length of data up to next full physical block */
81
82     length = (offset + hex_length + bsize - 1) & ~(bsize - 1);
83
84     buffer = (char *) malloc(length);
85     if (!buffer) {
86         fputs("alter: malloc failure!\n", stderr);
87         return;
88     }
89     if (ujfs_rw_diskblocks(fd, address, length, buffer, GET)) {
90         fputs("alter: failed reading disk data!\n", stderr);
91     }
92 }
```

```
91         free(buffer);
92         return;
93     }
94     for (ptr = hexstring; *ptr; ptr++) {
95         if (*ptr >= '0' && *ptr <= '9')
96             byte = *ptr - '0';
97         else if (*ptr >= 'a' && *ptr <= 'f')
98             byte = *ptr - 'a' + 10;
99         else if (*ptr >= 'A' && *ptr <= 'F')
100             byte = *ptr - 'A' + 10;
101         else {
102             fputs("alter: invalid hex digit!\n", stderr);
103             free(buffer);
104             return;
105         }
106
107         ptr++;
108         byte <<= 4;
109
110         if (*ptr >= '0' && *ptr <= '9')
111             byte += *ptr - '0';
112         else if (*ptr >= 'a' && *ptr <= 'f')
113             byte += *ptr - 'a' + 10;
114         else if (*ptr >= 'A' && *ptr <= 'F')
115             byte += *ptr - 'A' + 10;
116         else {
117             fputs("alter: invalid hex digit!\n", stderr);
118             free(buffer);
119             return;
120         }
121
122         buffer[offset++] = byte;
123     }
124     if (ujfs_rw_diskblocks(fd, address, length, buffer, PUT))
125         fputs("alter: failed writing disk data!\n", stderr);
126
127     free(buffer);
128     return;
129 }
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 /*
19 *   FUNCTION: Displays data in a variety of formats
20 */
21 #include "xpeek.h"
22 #include "jfs_endian.h"
23
24 #include <jfs_xtree.h>
25
26 /* libfs includes */
27 #include <devices.h>
28
29 static void display_hex(char *, unsigned, unsigned);
30
31 extern unsigned type_jfs;
32
33 void display()
34 {
35     int64_t addr;
36     int64_t block;
37     char *buffer;
38     char cmdline[512];
39     unsigned data_size;
40     int format = 'a';
41     int64_t len;
42     unsigned length = 0;
43     unsigned offset = 0;
44     char *token;
45
46     token = strtok(0, " \n");
47     if (token == 0) {
48         fputs("display: Please enter: block [offset [format [count]]]\ndisplay> ", stdout);
49         fgets(cmdline, 512, stdin);
50         token = strtok(cmdline, " \n");
51         if (token == 0)
52             return;
53     }
54     errno = 0;
55     block = strtoull(token, 0, 0);
56     if (block == 0 && errno) {
57         fputs("display: invalid block\n\n", stderr);
58         return;
59     }
60     if ((token = strtok(0, " \n"))) {
61         offset = strtoul(token, 0, 16);
62         if (offset == 0 && errno) {
63             fputs("display: invalid offset\n\n", stderr);
64             return;
65         }
66     }
67     if ((token = strtok(0, " \n")))
68         format = token[0];
69
70     if ((token = strtok(0, " \n"))) {
71         length = strtoul(token, 0, 0);
72         if (length == 0 && errno) {
73             fputs("display: invalid length\n\n", stderr);
74             return;
75         }
76     }
77
78     if (strtok(0, " \n")) {
79         fputs("display: Too many arguments\n\n", stderr);
80         return;
81     }
82
83     switch (format) {
84     case 'a':
85         data_size = 1;
86         if (length == 0)
87             length = bsize;
88         break;
89     case 'i':
90         data_size = sizeof (struct dinode);

```

```

91         if (length == 0)
92             length = 1;
93         break;
94     case 'I':
95         data_size = sizeof (struct iag);
96         if (length == 0)
97             length = 1;
98         break;
99     case 's':
100        data_size = sizeof (struct superbblock);
101        if (length == 0)
102            length = 1;
103        break;
104     case 'x':
105        data_size = 4;
106        if (length == 0)
107            length = bsize / 4;
108        break;
109     case 'X':
110        data_size = sizeof (xad_t);
111        if (length == 0)
112            length = 1;
113        break;
114     default:
115        fputs("display: invalid format\n\n", stderr);
116        return;
117     }
118
119     addr = block << l2bssize;
120     len = ((length * data_size) + offset + bsize - 1) & ~(bsize - 1);
121     buffer = malloc(len);
122     if (buffer == 0) {
123         fputs("display: error calling malloc\n\n", stderr);
124         return;
125     }
126
127     if (ujfs_rw_diskblocks(fd, addr, len, buffer, GET)) {
128         fputs("display: ujfs_rw_diskblocks failed\n\n", stderr);
129         free(buffer);
130         return;
131     }
132
133     printf("Block:%lld Real Address 0x%llx\n", (long long) block, (long long) addr);
134     switch (format) {
135     case 'a':
136     case 'x':
137         display_hex(&buffer[offset], length, offset);
138         break;
139     case 'i':
140         {
141             int i;
142             struct dinode *inode = (struct dinode *) &buffer[offset];
143             for (i = 0; i < length; i++, inode++) {
144                 /* swap if on big endian machine */
145                 ujfs_swap_dinode(inode, GET, type_jfs);
146                 display_inode(inode);
147                 if (more())
148                     return;
149             }
150         }
151         break;
152     case 'I':
153         {
154             int i;
155             struct iag *iag = (struct iag *) &buffer[offset];
156             for (i = 0; i < length; i++, iag++) {
157                 /* swap if on big endian machine */
158                 ujfs_swap_iag(iag);
159                 display_iag(iag);
160                 if (more())
161                     return;
162             }
163         }
164         break;
165     case 's':
166         /* swap if on big endian machine */
167         ujfs_swap_superblock((struct superbblock *) &buffer[offset]);
168         if (display_super((struct superbblock *) &buffer[offset]) == XPEEK_CHANGED) {
169             /* swap if on big endian machine */
170             ujfs_swap_superblock((struct superbblock *) &buffer[offset]);
171             if (ujfs_rw_diskblocks(fd, addr, len, buffer, PUT))
172                 fputs("Display: Error writing superbblock\n", stderr);
173         }
174         break;
175     default:
176         fputs("display: specified format not yet supported\n\n", stderr);
177         break;
178     }
179 }
180

```



```
181     free(buffer);
182     return;
183 }
184
185
186 /*
187  *   display_hex: display region in hex/ascii
188  */
189 static void display_hex(char *addr, unsigned length, unsigned offset)
190 {
191     uint8_t hextext[37];
192     uint8_t asciitxt[17];
193     uint8_t *x = (uint8_t *) addr, x1, x2;
194     int i, j, k, l;
195
196     hextext[36] = '\0';
197     asciitxt[16] = '\0';    /* null end of string */
198
199     l = 0;
200
201     for (i = 1; i <= ((length + 15) / 16); i++) {
202         if (i > 1 && ((i - 1) % 16) == 0)
203             if (more())
204                 break;
205
206         /* print address/offset */
207         printf("%08x:", offset + l);
208
209         /* print 16 bytes per line */
210         for (j = 0, k = 0; j < 16; j++, x++, l++) {
211             if ((j % 4) == 0)
212                 hextext[k++] = ' ';
213             if (l < length) {
214                 hextext[k++] = ((x1 = ((*x & 0xf0) >> 4) < 10)
215                     ? ('0' + x1) : ('A' + x1 - 10));
216                 hextext[k++] = ((x2 = (*x & 0x0f) < 10)
217                     ? ('0' + x2) : ('A' + x2 - 10));
218                 asciitxt[j] = ((*x < 0x20) || (*x >= 0x7f)) ? '.' : *x;
219             } else { /* byte not in range */
220                 hextext[k++] = ' ';
221                 hextext[k++] = ' ';
222                 asciitxt[j] = '.';
223             }
224         }
225         printf("%s |%s\n", hextext, asciitxt);
226     }
227 }
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 /*
19 *       dmap - display/modify disk map
20 */
21 #include "xpeek.h"
22
23 /* JFS includes */
24 #include <jfs_filsys.h>
25 #include <jfs_dmap.h>
26
27 /* libfs includes */
28 #include <jfs_endian.h>
29 #include <inode.h>
30 #include <devices.h>
31
32 /* ACTIONS */
33 #define DMAP_EXIT      0
34 #define DISPLAY_DBMAP  1
35 #define DISPLAY_CPAGE  2
36 #define DISPLAY_LEAF   3
37
38 /* DMAP BLOCK TYPES */
39 #define DMAP           -1
40 #define LEVEL0         0
41 #define LEVEL1         1
42 #define LEVEL2         2
43
44 /* DMAP BLOCK CALCULATIONS */
45 #define LOFACTOR        (LPERCTL + 1)
46 #define L1FACTOR        ((LOFACTOR * LPERCTL) + 1)
47 #define L1PAGE(l1)      (2 + ((l1) * L1FACTOR))
48 #define LOPAGE(l1, l0)  (L1PAGE(l1) + 1 + (l0 * LOFACTOR))
49 #define DMAPPAGE(l1, l0, d)  (LOPAGE(l1, l0) + 1 + d)
50
51 /* forward references */
52 int decode_pagenum(int64_t page, int *l1, int *l0, int *dmap);
53 int display_agfree(int64_t * agfree);
54 int display_cpagen(struct dinode *, int64_t *);
55 int display_dbmap(struct dbmap *, int64_t *, int64_t);
56 int display_leaf(struct dinode *, int64_t *);
57 int display_tree(struct dmapctl *, int64_t *, int *);
58
59 /* external references */
60 int display_map(unsigned *, int);      /* iag.c */
61
62 /* Global Data */
63 int dmap_level;          /* Maximum level of dtree */
64 int dmap_l2bpp;         /* log 2 of blocks per page */
65
66 extern unsigned type_jfs;
67
68 void dmap()
69 {
70     int action;
71     int64_t address;
72     struct dinode bmap_inode;
73     int64_t cntl_addr;
74     struct dbmap cntl_page;
75     int64_t lblock;      /* Logical block of BMAP to be displayed */
76
77     lblock = 0;
78     action = DISPLAY_DBMAP;
79
80     /* Read block allocation map Inode */
81     if (find_inode(BMAP_I, AGGREGATE_I, &address) ||
82         xRead(address, sizeof (struct dinode), (char *) &bmap_inode)) {
83         fputs("dmap: Error reading block allocation map inode\n", stderr);
84         return;
85     }
86
87     /* swap if on big endian machine */
88     ujfs_swap_dinode(&bmap_inode, GET, type_jfs);
89
90     /* Read overall control page for the map */

```

```

91
92     if (ujfs_rwdaddr(fd, &cntl_addr, &bmap_inode, (int64_t) 0, GET, bsize) ||
93         xRead(cntl_addr, sizeof (struct dbmap), (char *) &cntl_page)) {
94         fputs("dmap: Error reading aggregate dmap control page\n", stderr);
95         return;
96     }
97
98     /* swap if on big endian machine */
99     ujfs_swap_dbmap(&cntl_page);
100
101     /* Useful stuff */
102     dmap_level = BMAPSZTOLEV(cntl_page.dn_mapsize);
103     dmap_l2bpp = cntl_page.dn_l2nbperpage;
104
105     while (action != DMAP_EXIT) {
106         switch (action) {
107             case DISPLAY_DBMAP:
108                 action = display_dbmap(&cntl_page, &lblock, cntl_addr);
109                 break;
110             case DISPLAY_CPAGE:
111                 action = display_cpage(&bmap_inode, &lblock);
112                 break;
113             case DISPLAY_LEAF:
114                 action = display_leaf(&bmap_inode, &lblock);
115                 break;
116             default:
117                 fputs("dmap: Internal Error!\n", stderr);
118                 return;
119         }
120     }
121     return;
122 }
123
124 /*****
125 ***** Sample output from display_dbmap()
126
127 Block allocation map control page at block 128
128
129 [1] dn_mapsize:          0x00000031fa0   [9] dn_agheigth:          0
130 [2] dn_nfree:           0x00000030ed8   [10] dn_agwidth:           1
131 [3] dn_l2nbperpage:     3               [11] dn_agstart:           341
132 [4] dn_numag:           25              [12] dn_agl2size:          13
133 [5] dn_maxlevel:        0               [13] dn_agfree:            type 'f'
134 [6] dn_maxag:           16              [14] dn_agsize:            8192
135 [7] dn_aggpref:         16              [15] pad:                 Not Displayed
136 [8] dn_aglevel:         0
137
138 *****/
139
140 int display_dbmap(struct dbmap *hdr, int64_t *lblock, int64_t hdr_addr)
141 {
142     char cmdline[80];
143     int field;
144     int rc;
145     char *token;
146
147     dbmap_redisplay:
148     printf("\nBlock allocation map control page at block %lld\n",
149           (long long) (hdr_addr >> l2bsize));
150     printf("[1] dn_mapsize:\t\t0x%011llx\t", (long long) hdr->dn_mapsize);
151     printf("[9] dn_agheigth:\t%d\n", hdr->dn_agheigth);
152     printf("[2] dn_nfree:\t\t0x%011llx\t", (long long) hdr->dn_nfree);
153     printf("[10] dn_agwidth:\t%d\n", hdr->dn_agwidth);
154     printf("[3] dn_l2nbperpage:\t%d\t", hdr->dn_l2nbperpage);
155     printf("[11] dn_agstart:\t%d\n", hdr->dn_agstart);
156     printf("[4] dn_numag:\t\t%d\t", hdr->dn_numag);
157     printf("[12] dn_agl2size:\t%d\n", hdr->dn_agl2size);
158     printf("[5] dn_maxlevel:\t%d\t", hdr->dn_maxlevel);
159     printf("[13] dn_agfree:\t\ttype 'f'\n");
160     printf("[6] dn_maxag:\t\t%d\t", hdr->dn_maxag);
161     printf("[14] dn_agsize:\t\t%lld\n", (long long) hdr->dn_agsize);
162     printf("[7] dn_aggpref:\t\t%d\t", hdr->dn_aggpref);
163     printf("[15] pad:\t\tNot Displayed\n");
164     printf("[8] dn_aglevel:\t\t%d\n", hdr->dn_aglevel);
165
166     dbmap_retry:
167     fputs("display_dbmap: [m]odify, [f]ree count, [t]ree, e[x]it > ", stdout);
168     fgets(cmdline, 80, stdin);
169     token = strtok(cmdline, " \n");
170     if ((token == 0) || (token[0] == 't')) {
171         switch (dmap_level) {
172             case 2:
173                 *lblock = 1; /* L2 page */
174                 break;
175             case 1:
176                 *lblock = 2; /* L1 page */
177                 break;
178             case 0:
179                 *lblock = 3; /* L0 page */
180                 break;

```

```

181         default:
182             fprintf(stderr, "display_dbmap: dmap_level=%d\n", dmap_level);
183             return DMAP_EXIT;
184         }
185         return DISPLAY_CPAGE;
186     }
187     if (token[0] == 'f') {
188         rc = display_agfree(hdr->dn_agfree);
189
190         /* swap if on big endian machine */
191         ujfs_swap_dbmap(hdr);
192
193         if ((rc & XPEEK_CHANGED) && xWrite(hdr_addr, sizeof (struct dbmap), (char *) hdr)) {
194             fputs("display_dbmap: error writing control header\n", stderr);
195             /* swap if on big endian machine */
196             ujfs_swap_dbmap(hdr);
197             return DMAP_EXIT;
198         }
199
200         /* swap if on big endian machine */
201         ujfs_swap_dbmap(hdr);
202
203         if (rc & XPEEK_REDISPLAY)
204             goto dbmap_redisplay;
205
206         return DMAP_EXIT;
207     }
208     if (token[0] != 'm') { /* assume 'x' */
209         return DMAP_EXIT;
210     }
211
212     field = m_parse(cmdline, 14, &token);
213     if (field == 0)
214         goto dbmap_retry;
215
216     switch (field) {
217     case 1:
218         hdr->dn_mapsize = strtoull(token, 0, 16);
219         break;
220     case 2:
221         hdr->dn_nfree = strtoull(token, 0, 16);
222         break;
223     case 3:
224         hdr->dn_l2nbperpage = strtoul(token, 0, 0);
225         break;
226     case 4:
227         hdr->dn_numag = strtoul(token, 0, 0);
228         break;
229     case 5:
230         hdr->dn_maxlevel = strtoul(token, 0, 0);
231         break;
232     case 6:
233         hdr->dn_maxag = strtoul(token, 0, 0);
234         break;
235     case 7:
236         hdr->dn_agpref = strtoul(token, 0, 0);
237         break;
238     case 8:
239         hdr->dn_aglevel = strtoul(token, 0, 0);
240         break;
241     case 9:
242         hdr->dn_agheight = strtoul(token, 0, 0);
243         break;
244     case 10:
245         hdr->dn_agwidth = strtoul(token, 0, 0);
246         break;
247     case 11:
248         hdr->dn_agstart = strtoul(token, 0, 0);
249         break;
250     case 12:
251         hdr->dn_agl2size = strtoul(token, 0, 0);
252         break;
253     case 13:
254         fputs("display_dbmap: Can't change this field from here.\n", stderr);
255         goto dbmap_retry;
256     case 14:
257         hdr->dn_agsize = strtoull(token, 0, 0);
258         break;
259     }
260
261     /* swap if on big endian machine */
262     ujfs_swap_dbmap(hdr);
263
264     if (xWrite(hdr_addr, sizeof (struct dbmap), (char *) hdr)) {
265         fputs("display_dbmap: error writing control header\n", stderr);
266         /* swap back if on big endian machine */
267         ujfs_swap_dbmap(hdr);
268         return DMAP_EXIT;
269     }
270

```

```

271      /* swap back if on big endian machine */
272      ujfs_swap_dbmap(hdr);
273
274      goto dbmap_redisplay;
275  }
276
277  /******
278  ***** Example Output of display_cpage()
279
280  Level 0 Control Page at block 136
281
282  [1] nleafs:      1024          [5] budmin:      13
283  [2] l2nleafs:   10           [6] stree:       hit <enter>
284  [3] leafidx:    341          [7] pad:        Not Displayed
285  [4] height:     5
286
287  *****/
288  int display_cpage(struct dinode * bmap_inode, int64_t * lblock)
289  {
290      int64_t address;
291      int changed = 0;
292      char cmdline[80];
293      struct dmapctl ctl_page;
294      int field;
295      int rc;
296      char *token;
297      int dmap, l0, l1, type;
298
299      type = decode_pagenum(*lblock, &l1, &l0, &dmap);
300
301      if (ujfs_rwdaddr(fd, &address, bmap_inode, (*lblock) << dmap_l2bpp,
302          GET, bsize) || xRead(address, sizeof (struct dmapctl), (char *) &ctl_page)) {
303          fputs("display_cpage: Error reading control page!\n", stderr);
304          return DMAP_EXIT;
305      }
306
307      /* swap if on big endian machine */
308      ujfs_swap_dmapctl(&ctl_page);
309
310  cpage_redisplay:
311      printf("\nLevel %d Control Page at block %lld\n", type, (long long) (address >> l2bsize));
312
313      printf("[1] nleafs:%d\t\t", ctl_page.nleafs);
314      printf("[5] budmin:%d\n", ctl_page.budmin);
315      printf("[2] l2nleafs:%d\t\t", ctl_page.l2nleafs);
316      printf("[6] stree:%hit <enter>\n");
317      printf("[3] leafidx:%d\t\t", ctl_page.leafidx);
318      printf("[7] pad:%Not Displayed\n");
319      printf("[4] height:%d\n", ctl_page.height);
320
321  cpage_retry:
322      fputs("[m]odify, [u]p, [r]ight or [l]eft sibling, e[x]it, [s]tree > ", stdout);
323
324      fgets(cmdline, 80, stdin);
325      token = strtok(cmdline, " \n");
326      if ((token == 0) || (token[0] == 's')) {
327          rc = display_tree(&ctl_page, lblock, &changed);
328
329          /* swap if on big endian machine */
330          ujfs_swap_dmapctl(&ctl_page);
331
332          if (changed && xWrite(address, sizeof (struct dmapctl), (char *) &ctl_page)) {
333              fputs("display_cpage: Error writing control page!\n", stderr);
334              /* swap back if on big endian machine */
335              ujfs_swap_dmapctl(&ctl_page);
336              return DMAP_EXIT;
337          }
338
339          /* swap if on big endian machine */
340          ujfs_swap_dmapctl(&ctl_page);
341
342          return rc;
343      }
344      if (token[0] == 'u') {
345          if (type == LEVEL2) {
346              *lblock = 0;
347              return DISPLAY_DBMAP;
348          }
349          if (type == LEVEL1) {
350              if (dmap_level > 1) {
351                  *lblock = 1;
352                  return DISPLAY_CPAGE;
353              } else {
354                  *lblock = 0;
355                  return DISPLAY_DBMAP;
356              }
357          }
358          if (type == LEVEL0) {
359              if (dmap_level > 0) {
360                  *lblock = L1PAGE(l1);

```

```

361         return DISPLAY_CPAGE;
362     } else {
363         *lblock = 0;
364         return DISPLAY_DBMAP;
365     }
366 }
367
368 if ((token[0] == 'r') || (token[0] == 'l')) {
369     if (type == LEVEL2) {
370         fputs("Level 2 node has no siblings!\n", stderr);
371         goto cpage_retry;
372     }
373     if (type == LEVEL1) {
374         if (dmap_level < 2) {
375             fputs("Level 1 node has no siblings!\n", stderr);
376             goto cpage_retry;
377         }
378         if (token[0] == 'r') {
379             if (l1 == LPERCTL - 1) {
380                 fputs("No right sibling!\n", stderr);
381                 goto cpage_retry;
382             }
383             l1++;
384         } else {
385             if (l1 == 0) {
386                 fputs("No left sibling!\n", stderr);
387                 goto cpage_retry;
388             }
389             l1--;
390         }
391         *lblock = L1PAGE(l1);
392         return DISPLAY_CPAGE;
393     }
394     if (type == LEVEL0) {
395         if (dmap_level < 1) {
396             fputs("Level 0 node has no siblings!\n", stderr);
397             goto cpage_retry;
398         }
399         if (token[0] == 'r') {
400             if (l0 == LPERCTL - 1) {
401                 if ((dmap_level == 1) || (l1 == LPERCTL - 1)) {
402                     fputs("No right sibling!\n", stderr);
403                     goto cpage_retry;
404                 }
405                 l1++;
406                 l0 = 0;
407             } else
408                 l0++;
409         } else {
410             if (l0 == 0) {
411                 if (l1 == 0) {
412                     fputs("No left sibling!\n", stderr);
413                     goto cpage_retry;
414                 }
415                 l1--;
416                 l0 = LPERCTL - 1;
417             } else
418                 l0--;
419         }
420         *lblock = L0PAGE(l1, l0);
421         return DISPLAY_CPAGE;
422     }
423     fprintf(stderr, "display_cpage: decode_pagenum returned type %d\n", type);
424     return DMAP_EXIT;
425 }
426 if (token[0] == 'm') {
427     field = m_parse(cmdline, 5, &token);
428     if (field == 0)
429         goto cpage_retry;
430
431     switch (field) {
432     case 1:
433         ctl_page.nleafs = strtoul(token, 0, 0);
434         break;
435     case 2:
436         ctl_page.l2nleafs = strtoul(token, 0, 0);
437         break;
438     case 3:
439         ctl_page.leafidx = strtoul(token, 0, 0);
440         break;
441     case 4:
442         ctl_page.height = strtoul(token, 0, 0);
443         break;
444     case 5:
445         ctl_page.budmin = strtoul(token, 0, 0);
446         break;
447     }
448
449     /* swap if on big endian machine */
450     ujfs_swap_dmapctl(&ctl_page);

```

```

451         if (xWrite(address, sizeof (struct dmapctl), (char *) &ctl_page)) {
452             fputs("display_cpage: Error writing control page!\n", stderr);
453             /* swap back if on big endian machine */
454             ujfs_swap_dmapctl(&ctl_page);
455             return DMAP_EXIT;
456         }
457     }
458
459     /* swap back if on big endian machine */
460     ujfs_swap_dmapctl(&ctl_page);
461
462     goto cpage_redisplay;
463 }
464 return DMAP_EXIT;
465 }
466
467 /*****
468 ***** Example output of display_leaf()
469
470 Dmap Page at block 144
471
472 [1] nblocks:          8192          [8] tree.budmin:          5
473 [2] nfree:           7624          [9] tree.stree:           Hit <enter>
474 [3] start:           0             [10] tree.pad:           Not Displayed
475 [4] tree.nleafs:     256           [11] pad:               Not Displayed
476 [5] tree.l2nleafs:   8             [12] wmap:              Type 'w'
477 [6] tree.leafidx:    85           [13] pmap:              Type 'p'
478 [7] tree.height:     4
479
480 *****/
481
482 int display_leaf(struct dinode * bmap_inode, int64_t * lblock)
483 {
484     int64_t address;
485     int changed = 0;
486     char cmdline[80];
487     struct dmap d_map;
488     int field;
489     int rc;
490     char *token;
491     int dmap, l0, l1, type;
492
493     type = decode_pagenum(*lblock, &l1, &l0, &dmap);
494
495     if (ujfs_rwdaddr(fd, &address, bmap_inode, (*lblock) << dmap_l2bpp,
496         GET, bsize) || xRead(address, sizeof (struct dmap), (char *) &d_map)) {
497         fputs("display_leaf: Error reading dmap page!\n", stderr);
498         return DMAP_EXIT;
499     }
500
501     /* swap if on big endian machine */
502     ujfs_swap_dmap(&d_map);
503
504     leaf_redisplay:
505     printf("\nDmap Page at block %lld\n", (long long) (address >> l2bsize));
506
507     printf("[1] nblocks:\t\t%d\t\t", d_map.nblocks);
508     printf("[8] tree.budmin:\t\t%d\n", d_map.tree.budmin);
509     printf("[2] nfree:\t\t\t%d\t\t", d_map.nfree);
510     printf("[9] tree.stree:\t\tHit<enter>\n");
511     printf("[3] start:\t\t\t%lld\t\t", (long long) d_map.start);
512     printf("[10] tree.pad:\t\tNot Displayed\n");
513     printf("[4] tree.nleafs:\t\t%d\t\t", d_map.tree.nleafs);
514     printf("[11] pad:\t\t\tNot Displayed\n");
515     printf("[5] tree.l2nleafs:\t\t%d\t\t", d_map.tree.l2nleafs);
516     printf("[12] wmap:\t\t\tType 'w'\n");
517     printf("[6] tree.leafidx:\t\t%d\t\t", d_map.tree.leafidx);
518     printf("[13] pmap:\t\t\tType 'p'\n");
519     printf("[7] tree.height:\t\t%d\n", d_map.tree.height);
520
521     leaf_retry:
522     fputs("[m]odify, [u]p, [r]ight or [l]eft, [w]map, [p]map, e[x]it, [s]tree > ", stdout);
523     fgets(cmdline, 80, stdin);
524     token = strtok(cmdline, " \n");
525
526     if ((token == 0) || (token[0] == 's')) {
527         rc = display_tree((struct dmapctl *) &d_map.tree, lblock, &changed);
528         /* swap if on big endian machine */
529         ujfs_swap_dmap(&d_map);
530
531         if (changed && xWrite(address, sizeof (struct dmap), (char *) &d_map)) {
532             fputs("display_leaf: Error writing dmap page!\n", stderr);
533             /* swap back if on big endian machine */
534             ujfs_swap_dmap(&d_map);
535             return DMAP_EXIT;
536         }
537
538         /* swap back if on big endian machine */
539         ujfs_swap_dmap(&d_map);
540

```

```

541         return rc;
542     }
543     if (token[0] == 'u') {
544         *lblock = LOPAGE(l1, 10);
545         return DISPLAY_CPAGE;
546     }
547     if (token[0] == 'r') {
548         if (dmap < LPERCTL - 1) {
549             *lblock = DMAPPAGE(l1, 10, dmap + 1);
550             return (DISPLAY_LEAF);
551         }
552         if (dmap_level > 0) {
553             if (l0 < LPERCTL - 1) {
554                 *lblock = DMAPPAGE(l1, 10 + 1, 0);
555                 return (DISPLAY_LEAF);
556             }
557             if ((dmap_level == 2) && (l1 < LPERCTL - 1)) {
558                 *lblock = DMAPPAGE(l1 + 1, 0, 0);
559                 return (DISPLAY_LEAF);
560             }
561         }
562         fputs("display_leaf: No right sibling.\n", stderr);
563         goto leaf_retry;
564     }
565     if (token[0] == 'l') {
566         if (dmap > 0) {
567             *lblock = DMAPPAGE(l1, 10, dmap - 1);
568             return (DISPLAY_LEAF);
569         }
570         if (dmap_level > 0) {
571             if (l0 > 0) {
572                 *lblock = DMAPPAGE(l1, 10 - 1, LPERCTL - 1);
573                 return (DISPLAY_LEAF);
574             }
575             if ((dmap_level == 2) && (l1 > 0)) {
576                 *lblock = DMAPPAGE(l1 - 1, LPERCTL - 1, LPERCTL - 1);
577                 return (DISPLAY_LEAF);
578             }
579         }
580         fputs("display_leaf: No left sibling.\n", stderr);
581         goto leaf_retry;
582     }
583     if ((token[0] == 'p') || (token[0] == 'w')) {
584         if (token[0] == 'p')
585             rc = display_map(d_map.pmap, LPERDMAP);
586         else
587             rc = display_map(d_map.wmap, LPERDMAP);
588
589         /* swap if on big endian machine */
590         ujfs_swap_dmap(&d_map);
591
592         if ((rc & XPEEK_CHANGED) && xWrite(address, sizeof(struct dmap), (char *) &d_map)) {
593             fputs("display_leaf: Error writing dmap page!\n", stderr);
594             /* swap back if on big endian machine */
595             ujfs_swap_dmap(&d_map);
596             return DMAP_EXIT;
597         }
598
599         /* swap back if on big endian machine */
600         ujfs_swap_dmap(&d_map);
601
602         if (rc & XPEEK_REDISPLAY)
603             goto leaf_redisplay;
604         return DMAP_EXIT;
605     }
606     if (token[0] != 'm') /* assume 'x' */
607         return DMAP_EXIT;
608
609     field = m_parse(cmdline, 8, &token);
610     if (field == 0)
611         goto leaf_retry;
612
613     switch (field) {
614     case 1:
615         d_map.nblocks = strtoul(token, 0, 0);
616         break;
617     case 2:
618         d_map.nfree = strtoul(token, 0, 0);
619         break;
620     case 3:
621         d_map.start = strtoull(token, 0, 0);
622         break;
623     case 4:
624         d_map.tree.nleafs = strtoul(token, 0, 0);
625         break;
626     case 5:
627         d_map.tree.l2nleafs = strtoul(token, 0, 0);
628         break;
629     case 6:
630         d_map.tree.leafidx = strtoul(token, 0, 0);

```



```

721     fgets(cmdline, 80, stdin);
722     token = strtok(cmdline, "          \n");
723     if ((token == 0) || (token[0] == 'x'))
724         return DMAP_EXIT;
725     if (token[0] == 'l') {
726         if (top_index == 0) { /* Need to move to another page */
727             fputs("For now, go [b]ack, then go left\n", stdout);
728             goto tree_retry;
729         }
730         top_index--;
731         goto redisplay_tree;
732     }
733     if (token[0] == 'r') {
734         if (top_index == (1 << ((ctlpage->height - top_level) << 1)) - 1) {
735             fputs("For now, [b]ack then go right\n", stdout);
736             goto tree_retry;
737         }
738         top_index++;
739         goto redisplay_tree;
740     }
741     if (token[0] == 'g') {
742         token = strtok(0, "          \n");
743         if (token == 0) {
744             fputs("Please enter: level index>", stdout);
745             fgets(cmdline, 80, stdin);
746             token = strtok(cmdline, "          \n");
747             if (token == 0)
748                 goto tree_retry;
749         }
750         level = strtoul(token, 0, 0);
751         token = strtok(0, "          \n");
752         if (token == 0) {
753             fputs("Not enough arguments\n", stderr);
754             goto tree_retry;
755         }
756         index = strtoul(token, 0, 0);
757         if ((level < 0) || (level > ctlpage->height) || (index < 0) ||
758             (index >= (1 << ((ctlpage->height - level) << 1)))) {
759             fputs("Invalid level and/or index\n", stderr);
760             goto tree_retry;
761         }
762         if ((level == 1) && ctlpage->height > 1) {
763             level++;
764             index >= 2;
765         } else if (level == 0) {
766             if (ctlpage->height == 1) { /* is this possible? */
767                 level++;
768                 index >= 2;
769             } else {
770                 level += 2;
771                 index >= 4;
772             }
773         }
774         top_level = level;
775         top_index = index;
776         goto redisplay_tree;
777     }
778     if (token[0] == 'u') {
779         if (top_level == ctlpage->height) { /* At top of tree */
780             fputs("Already at top of tree.\n", stdout);
781             goto tree_retry;
782         }
783         /* Move up one level */
784         top_level++;
785         top_index >= 2;
786         goto redisplay_tree;
787     }
788     if (token[0] == 'd') {
789         if (type == DMAP) {
790             fputs("[d]escend only valid for control pages\n", stderr);
791             goto tree_retry;
792         }
793         token = strtok(0, "          \n");
794         if (token == 0) {
795             fputs("Please enter: leaf#>", stdout);
796             fgets(cmdline, 80, stdin);
797             token = strtok(cmdline, "          \n");
798             if (token == 0)
799                 goto tree_retry;
800         }
801         index = strtoul(token, 0, 0);
802         if ((index < 0) || (index >= ctlpage->nleafs)) {
803             fputs("Invalid leaf index\n", stderr);
804             goto tree_retry;
805         }
806         if (type == LEVEL0) {
807             *lblock = DMAPPAGE(11, 10, index);
808             return DISPLAY_LEAF;
809         }
810         if (type == LEVEL1)

```

```

811         *lblock = LOPAGE(l1, index);
812         else /* LEVEL2 */
813             *lblock = L1PAGE(index);
814
815         return DISPLAY_CPAGE;
816     }
817     if (token[0] != 'm') { /* Assuming 'b' */
818         if (type == DMAP)
819             return DISPLAY_LEAF;
820         else
821             return DISPLAY_CPAGE;
822     }
823     token = strtok(0, " \n");
824     if (token == 0) {
825         fputs("Please enter: level index value >", stdout);
826         fgets(cmdline, 80, stdin);
827         token = strtok(0, " \n");
828         if (token == 0)
829             goto tree_retry;
830     }
831     level = strtol(token, 0, 0);
832     token = strtok(0, " \n");
833     if (token == 0) {
834         fputs("Not enough arguments!\n", stderr);
835         goto tree_retry;
836     }
837     index = strtol(token, 0, 0);
838     token = strtok(0, " \n");
839     if (token == 0) {
840         fputs("Not enough arguments!\n", stderr);
841         goto tree_retry;
842     }
843     if ((level < 0) || (level > ctlpage->height) || (index < 0) ||
844         (index >= (1 << ((ctlpage->height - level) << 1)))) {
845         fputs("Invalid level and/or index\n", stderr);
846         goto tree_retry;
847     }
848     ctlpage->stree[tree_offset[ctlpage->height - level] + index] = strtol(token, 0, 0);
849     *changed = 1;
850     goto redisplay_tree;
851 }
852
853 int decode_pagenum(int64_t page, int *l1, int *l0, int *dmap)
854 {
855     int remainder;
856
857     if (page == 0)
858         return -1;
859
860     if (page == 1)
861         return LEVEL2;
862
863     *l1 = (page - 2) / L1FACTOR;
864     remainder = (page - 2) % L1FACTOR;
865     if (remainder == 0)
866         return LEVEL1;
867
868     *l0 = (remainder - 1) / L0FACTOR;
869     remainder = (remainder - 1) % L0FACTOR;
870     if (remainder == 0)
871         return LEVEL0;
872
873     *dmap = remainder - 1;
874     return DMAP;
875 }
876
877 int display_agfree(int64_t * agfree)
878 {
879     char cmdline[80];
880     int end;
881     int i;
882     int index;
883     int rc = XPEEK_OK;
884     int start = 0;
885     char *token;
886
887     agfree_display:
888     end = MIN(start + 64, MAXAG);
889     for (i = start; i < end; i += 4)
890         printf("%3d 0x%016llx 0x%016llx 0x%016llx 0x%016llx\n", i, (long long) agfree[i],
891             (long long) agfree[i + 1], (long long) agfree[i + 2],
892             (long long) agfree[i + 3]);
893     agfree_retry:
894     fputs("display_agfree: [m]odify, [b]ack, e[x]it > ", stdout);
895     fgets(cmdline, 80, stdin);
896     token = strtok(cmdline, " \n");
897     if (token == 0) {
898         start = (end < MAXAG) ? end : 0;
899         goto agfree_display;
900     }

```

```
901     if (token[0] == 'x')
902         return rc;
903     if (token[0] != 'm') { /* assuming 'b' */
904         return (rc | XPEEK_REDISPLAY);
905     }
906
907     index = m_parse(cmdline, MAXAG - 1, &token);
908     if (index == 0)
909         goto agfree_retry;
910
911     agfree[index] = strtoull(token, 0, 16);
912     rc = XPEEK_CHANGED;
913     goto agfree_display;
914 }
```

```

1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 /*
19 *   FUNCTION: Explain how to use the xpeek utility
20 */
21 #include "xpeek.h"
22
23 void help()
24 {
25     char *cmd;
26     int cmd_len;
27
28     cmd = strtok(0, "\n");          /* space & tab */
29     if (strtok(0, "\n")) {
30         fputs("help: called with too many arguments\n", stderr);
31         return;
32     }
33     if (cmd == 0) {
34         fputs("\t\tXpeek Commands\n", stdout);
35         fputs("a[iter] <block> <offset> <hex string>\n", stdout);
36         fputs("b[tree] <block> [<offset>]\n", stdout);
37         fputs("cb[blfsc] \n", stdout);
38         fputs("dir[ectory] <inode number> [<fileset>]\n", stdout);
39         fputs("d[isplay] [<block> [<offset> [<format> [<count>]]]\n", stdout);
40         fputs("dm[ap] \n", stdout);
41         fputs("dt[ree] <inode number> [<fileset>]\n", stdout);
42         fputs("fscw[sphdr] \n", stdout);
43         fputs("h[elp] [<command>]\n", stdout);
44         fputs("ia[g] [<IAG number>] [a|s|<fileset>]\n", stdout);
45         fputs("i[node] [<inode number>] [a|s|<fileset>]\n", stdout);
46         fputs("logs[uper] \n", stdout);
47         fputs("q[uit] \n", stdout);
48         fputs("se[t] [<variable> <value>]\n", stdout);
49         fputs("su[perblock] [p|s] \n", stdout);
50         fputs("s2p[erblock] [p|s] \n", stdout);
51         fputs("u[nset] <variable> \n", stdout);
52         fputs("xt[ree] <inode number> [<fileset>]\n", stdout);
53         return;
54     }
55     cmd_len = strlen(cmd);
56     if (strncmp(cmd, "alter", cmd_len) == 0) {
57         fputs("a[iter] <block> <offset> <hex string>\n", stdout);
58         fputs("\t<block>\t\tblock number (decimal)\n", stdout);
59         fputs("\t<offset>\t\toffset within block (hex)\n", stdout);
60         fputs("\t<hex string>\t\tstring of hex digits\n", stdout);
61         fputs("alters disk data. <hex string> should contain an even number of digits\n",
62             stdout);
63     } else if (strncmp(cmd, "btree", cmd_len) == 0) {
64         fputs("b[tree] <block> [<offset>]\n", stdout);
65         fputs("\t<block>\t\tblock number (decimal)\n", stdout);
66         fputs("\t<offset>\t\toffset within block (hex)\n", stdout);
67         fputs("displays one node of a btree and enters a subcommand mode in which to\n",
68             stdout);
69         fputs("navigate the btree. Subcommands:\n", stdout);
70         fputs("\t\tvisit left sibling\n", stdout);
71         fputs("\t\tmodify current node\n", stdout);
72         fputs("\t\tvisit parent node\n", stdout);
73         fputs("\t\tvisit right sibling\n", stdout);
74         fputs("\t[0-9]\t\tvisit the nth child node\n", stdout);
75         fputs("\t\t\texit subcommand mode\n", stdout);
76     } else if (cmd_len > 1 && strncmp(cmd, "cbbfsc", cmd_len) == 0) {
77         fputs("cb[blfsc] \n", stdout);
78         fputs("Displays the area used by ClearBadBlockList \n", stdout);
79         fputs("for communication with fscd.\n", stdout);
80     } else if (cmd_len > 2 && strncmp(cmd, "directory", cmd_len) == 0) {
81         fputs("dir[ectory] <inode number> [<fileset>]\n", stdout);
82         fputs("\t<inode number>\t\tinode number of directory (decimal)\n", stdout);
83         fputs("\t<fileset>\t\tfileset number, currently must be zero\n", stdout);
84         fputs("Displays directory entries. Subcommands\n", stdout);
85         fputs("\t\t\tmodify entries\n", stdout);
86         fputs("\t\t\texit subcommand mode\n", stdout);
87     } else if (cmd_len > 1 && strncmp(cmd, "dtree", cmd_len) == 0) {
88         fputs("dt[ree] <inode number> [<fileset>]\n", stdout);
89         fputs("\t<inode number>\t\tinode number of directory (decimal)\n", stdout);
90         fputs("\t<fileset>\t\tfileset number, currently must be zero\n", stdout);

```



```
181         ("Displays root of the non-directory btree and enters a subcommand mode in which to\n",
182          stdout);
183         fputs("navigate the btree. Subcommands:\n", stdout);
184         fputs("\tl\tvisit left sibling\n", stdout);
185         fputs("\tm\tmodify current node\n", stdout);
186         fputs("\tp\tvisit parent node (not parent directory)\n", stdout);
187         fputs("\tr\tvisit right sibling\n", stdout);
188         fputs("\t[0-9]\tvisit the nth child node\n", stdout);
189         fputs("\tx\texit subcommand mode\n\n", stdout);
190     } else
191         fprintf(stderr, "help: Unknown command '%s'\n\n", cmd);
192 }
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 /*
19 *   io.c - I/O routines
20 */
21 #include "xpeek.h"
22
23 /* libfs includes */
24 #include <devices.h>
25
26 int xRead(int64_t address, unsigned count, char *buffer)
27 {
28     int64_t block_address;
29     char *block_buffer;
30     int64_t length;
31     unsigned offset;
32
33     offset = address & (bsize - 1);
34     length = (offset + count + bsize - 1) & ~(bsize - 1);
35
36     if ((offset == 0) & (length == count))
37         return ujfs_rw_diskblocks(fd, address, count, buffer, GET);
38
39     block_address = address - offset;
40     block_buffer = (char *) malloc(length);
41     if (block_buffer == 0)
42         return 1;
43
44     if (ujfs_rw_diskblocks(fd, block_address, length, block_buffer, GET)) {
45         free(block_buffer);
46         return 1;
47     }
48     memcpy(buffer, block_buffer + offset, count);
49     free(block_buffer);
50     return 0;
51 }
52
53 int xWrite(int64_t address, unsigned count, char *buffer)
54 {
55     int64_t block_address;
56     char *block_buffer;
57     int64_t length;
58     unsigned offset;
59
60     offset = address & (bsize - 1);
61     length = (offset + count + bsize - 1) & ~(bsize - 1);
62
63     if ((offset == 0) & (length == count))
64         return ujfs_rw_diskblocks(fd, address, count, buffer, PUT);
65
66     block_address = address - offset;
67     block_buffer = (char *) malloc(length);
68     if (block_buffer == 0)
69         return 1;
70
71     if (ujfs_rw_diskblocks(fd, block_address, length, block_buffer, GET)) {
72         free(block_buffer);
73         return 1;
74     }
75     memcpy(block_buffer + offset, buffer, count);
76     if (ujfs_rw_diskblocks(fd, block_address, length, block_buffer, PUT)) {
77         free(block_buffer);
78         return 1;
79     }
80     free(block_buffer);
81     return 0;
82 }
```



```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 /*
19 *   ui.c - User Interface routines
20 */
21 #include "xpeek.h"
22
23 /*
24 *   m_parse - parse parameters to 'm'odify subcommand
25 *
26 *   NOTE: Assumes last call to strtok() parsed "m" subcommand from
27 *   command line.
28 */
29 int m_parse(char *cmd_line, int n_fields, char **value)
30 {
31     int field_number;
32     char *token;
33
34     token = strtok(0, " \n");
35     if (token == 0) {
36         fputs("Please enter: field-number value> ", stdout);
37         fgets(cmd_line, 80, stdin);
38         token = strtok(cmd_line, " \n");
39         if (token == 0)
40             return 0;
41     }
42     field_number = strtol(token, 0, 0);
43     if (field_number < 1 || field_number > n_fields) {
44         fputs("Invalid field number\n", stderr);
45         return 0;
46     }
47     *value = strtok(0, " \n");
48     if (*value == 0) {
49         fputs("Not enough arguments\n", stderr);
50         return 0;
51     }
52     if (strtok(0, " \n")) {
53         fputs("Too many arguments\n", stderr);
54         return 0;
55     }
56     return field_number;
57 }
```

```
1  /*
2  *   Copyright (c) International Business Machines Corp., 2000-2002
3  *
4  *   This program is free software; you can redistribute it and/or modify
5  *   it under the terms of the GNU General Public License as published by
6  *   the Free Software Foundation; either version 2 of the License, or
7  *   (at your option) any later version.
8  *
9  *   This program is distributed in the hope that it will be useful,
10 *   but WITHOUT ANY WARRANTY; without even the implied warranty of
11 *   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See
12 *   the GNU General Public License for more details.
13 *
14 *   You should have received a copy of the GNU General Public License
15 *   along with this program; if not, write to the Free Software
16 *   Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
17 */
18 /*
19 *   FUNCTION: common data & function prototypes
20 */
21 #include <jfs_types.h>
22
23 /* system includes */
24 #include <stdio.h>
25 #include <string.h>
26 #include <stdlib.h>
27 #include <errno.h>
28
29 /* JFS includes */
30 #include <jfs_dinode.h>
31 #include <jfs_imap.h>
32 #include <jfs_superblock.h>
33 #include <devices.h>
34
35 /* Defines */
36 #define AGGREGATE_2ND_I -1
37
38 #define XPEEK_OK        0x00
39 #define XPEEK_CHANGED  0x01
40 #define XPEEK_REDISPLAY 0x10
41 #define XPEEK_ERROR    -1
42
43 /* Global Data */
44 extern int bsize;
45 extern HFILE fd;
46 extern short l2bsize;
47
48 /* xpeek functions */
49
50 void alter(void);
51 void cbbfscck(void);
52 void directory(void);
53 void display(void);
54 void display_iag(struct iag *);
55 void display_inode(struct dinode *);
56 int display_super(struct superblock *);
57 void dmap(void);
58 void dtree(void);
59 void help(void);
60 int find_iag(unsigned iagnum, unsigned which_table, int64_t * address);
61 int find_inode(unsigned inum, unsigned which_table, int64_t * address);
62 void fsckwspchr(void);
63 void iag(void);
64 void inode(void);
65 void logsuper(void);
66 int m_parse(char *, int, char **);
67 int more(void);
68 char prompt(char *);
69 void superblock(void);
70 void s2perblock(void);
71 void xtree(void);
72
73 int xRead(int64_t, unsigned, char *);
74 int xWrite(int64_t, unsigned, char *);
75
76 #define fputs(string,fd) { fputs(string,fd); fflush(fd); }
```

1	./JFS/jfsutils/fsck/fsckbmap.c	Pages	1- 23	2046 lines
2	./JFS/jfsutils/fsck/fsckconn.c	Pages	24- 37	1249 lines
3	./JFS/jfsutils/fsck/fsckpfs.h	Pages	38- 38	25 lines
4	./JFS/jfsutils/fsck/xchkdsk.h	Pages	39- 39	40 lines
5	./JFS/jfsutils/fsck/xfck.h	Pages	40- 47	672 lines
6	./JFS/jfsutils/fsck/xfckint.h	Pages	48- 52	404 lines
7	./JFS/jfsutils/include/jfs_btree.h	Pages	53- 53	50 lines
8	./JFS/jfsutils/include/jfs_dinode.h	Pages	54- 55	170 lines
9	./JFS/jfsutils/include/jfs_dmap.h	Pages	56- 58	218 lines
10	./JFS/jfsutils/include/jfs_dtree.h	Pages	59- 61	231 lines
11	./JFS/jfsutils/include/jfs_filsys.h	Pages	62- 64	260 lines
12	./JFS/jfsutils/include/jfs_imap.h	Pages	65- 66	117 lines
13	./JFS/jfsutils/include/jfs_superblock.h	Pages	67- 68	114 lines
14	./JFS/jfsutils/include/jfs_xtree.h	Pages	69- 70	101 lines
15	./JFS/jfsutils/libfs/debug.h	Pages	71- 71	43 lines
16	./JFS/jfsutils/libfs/diskmap.c	Pages	72- 75	314 lines
17	./JFS/jfsutils/libfs/diskmap.h	Pages	76- 76	30 lines
18	./JFS/jfsutils/libfs/fsck_base.h	Pages	77- 77	50 lines
19	./JFS/jfsutils/libfs/fsckcdbl.h	Pages	78- 78	71 lines
20	./JFS/jfsutils/libfs/fscklog.h	Pages	79- 79	73 lines
21	./JFS/jfsutils/libfs/fsckwsp.h	Pages	80- 89	866 lines
22	./JFS/jfsutils/libfs/inode.h	Pages	90- 90	28 lines
23	./JFS/jfsutils/libfs/jfs_endian.h	Pages	91- 91	84 lines
24	./JFS/jfsutils/libfs/libjufs.h	Pages	92- 92	26 lines
25	./JFS/jfsutils/libfs/logform.h	Pages	93- 93	24 lines
26	./JFS/jfsutils/libfs/utlsub.c	Pages	94- 94	82 lines
27	./JFS/jfsutils/libfs/utlsub.h	Pages	95- 95	33 lines
28	./JFS/jfsutils/logdump/helpers.c	Pages	96- 97	115 lines
29	./JFS/jfsutils/mkfs/initmap.h	Pages	98- 98	32 lines
30	./JFS/jfsutils/mkfs/inodemap.h	Pages	99- 99	26 lines
31	./JFS/jfsutils/tune/super.c	Pages	100-102	245 lines
32	./JFS/jfsutils/tune/tune.c	Pages	103-107	430 lines
33	./JFS/jfsutils/xpeek/alter.c	Pages	108-109	130 lines
34	./JFS/jfsutils/xpeek/display.c	Pages	110-112	228 lines
35	./JFS/jfsutils/xpeek/dmap.c	Pages	113-123	915 lines
36	./JFS/jfsutils/xpeek/help.c	Pages	124-126	193 lines
37	./JFS/jfsutils/xpeek/io.c	Pages	127-127	83 lines
38	./JFS/jfsutils/xpeek/ui.c	Pages	128-128	58 lines
39	./JFS/jfsutils/xpeek/xpeek.h	Pages	129-129	77 lines

End of Table of Contents