

```

*****
61235 Thu Oct 23 11:04:51 2014
new/usr/src/uts/common/avs/ns/rdc/rdc_bitmap.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
unchanged portion omitted

620 static int
621 rdc_read_bitmap(rdc_k_info_t *krdc, struct bitmapdata *data)
622 {
623     rdc_u_info_t *urdc;
624     int sts;

626     if (krdc == NULL) {
627         return (-1);
628     }

630     if (data != NULL) {
631         data->data = kmem_alloc(krdc->bitmap_size, KM_SLEEP);
632         data->len = krdc->bitmap_size;

634         if (data->data == NULL) {
635             return (-1);
636         }
633     }

635     mutex_enter(&krdc->bmapmutex);

637     urdc = &rdc_u_info[krdc->index];
638     if (rdc_get_vflags(urdc) & RDC_BMP_FAILED) {
639         mutex_exit(&krdc->bmapmutex);
640         return (-1);
641     }

643     if (krdc->bitmapfd == NULL) {
644         mutex_exit(&krdc->bmapmutex);
645         return (-1);
646     }

648     if (data == NULL && krdc->dcio_bitmap == NULL) {
649         mutex_exit(&krdc->bmapmutex);
650         return (-1);
651     }

653     if (_rdc_rsrv_devs(krdc, RDC_BMP, RDC_INTERNAL)) {
654         cmn_err(CE_WARN, "!rdc_read_bitmap: %s reserve failed",
655             urdc->primary.file);
656         rdc_set_flags_log(urdc, RDC_BMP_FAILED, "reserve failed");
657         mutex_exit(&krdc->bmapmutex);
658         return (-1);
659     }

661     if (krdc->bmp_kstats) {
662         mutex_enter(krdc->bmp_kstats->ks_lock);
663         kstat_runq_enter(KSTAT_IO_PTR(krdc->bmp_kstats));
664         mutex_exit(krdc->bmp_kstats->ks_lock);
665     }

667     sts = rdc_ns_io(krdc->bitmapfd, NSC_RDBUF, RDC_BITMAP_FBA,
668         data ? data->data : krdc->dcio_bitmap, krdc->bitmap_size);

670     if (krdc->bmp_kstats) {
671         mutex_enter(krdc->bmp_kstats->ks_lock);
672         kstat_runq_exit(KSTAT_IO_PTR(krdc->bmp_kstats));
673         mutex_exit(krdc->bmp_kstats->ks_lock);

```

```

674         KSTAT_IO_PTR(krdc->bmp_kstats)->reads++;
675         KSTAT_IO_PTR(krdc->bmp_kstats)->nread += krdc->bitmap_size;
676     }

678     _rdc_rlse_devs(krdc, RDC_BMP);

680     if (!RDC_SUCCESS(sts)) {
681         cmn_err(CE_WARN, "!rdc_read_bitmap: %s read failed",
682             urdc->primary.file);
683         rdc_set_flags_log(urdc, RDC_BMP_FAILED, "read failed");
684         mutex_exit(&krdc->bmapmutex);
685         return (-1);
686     }

688     mutex_exit(&krdc->bmapmutex);
689     return (0);
690 }
unchanged portion omitted

```

```

*****
160953 Thu Oct 23 11:04:51 2014
new/usr/src/uts/common/avs/ns/rdc/rdc_io.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

305 /*
306 * _rdc_load() - rdc is being loaded, Allocate anything
307 * that will be needed while the cache is loaded but doesn't really
308 * depend on configuration parameters.
309 *
310 */
311 int
312 _rdc_load(void)
313 {
314     int i;
315     rdc_k_info_t *krdc;

317     mutex_init(&rdc_ping_lock, NULL, MUTEX_DRIVER, NULL);
318     mutex_init(&net_blk_lock, NULL, MUTEX_DRIVER, NULL);
319     mutex_init(&rdc_conf_lock, NULL, MUTEX_DRIVER, NULL);
320     mutex_init(&rdc_many_lock, NULL, MUTEX_DRIVER, NULL);
321     mutex_init(&rdc_net_hnd_id_lock, NULL, MUTEX_DRIVER, NULL);
322     mutex_init(&rdc_clnt_lock, NULL, MUTEX_DRIVER, NULL);
323     mutex_init(&sync_info.lock, NULL, MUTEX_DRIVER, NULL);

325 #ifdef DEBUG
326     mutex_init(&rdc_cntlock, NULL, MUTEX_DRIVER, NULL);
327 #endif

329     if ((i = nsc_max_devices()) < rdc_max_sets)
330         rdc_max_sets = i;
331     /* following case for partial installs that may fail */
332     if (!rdc_max_sets)
333         rdc_max_sets = 1024;

335     rdc_k_info = kmem_zalloc(sizeof (*rdc_k_info) * rdc_max_sets, KM_SLEEP);
336     if (!rdc_k_info)
337         return (ENOMEM);

337     rdc_u_info = kmem_zalloc(sizeof (*rdc_u_info) * rdc_max_sets, KM_SLEEP);
338     if (!rdc_u_info) {
339         kmem_free(rdc_k_info, sizeof (*rdc_k_info) * rdc_max_sets);
340         return (ENOMEM);
341     }

339     net_exit = ATM_NONE;
340     for (i = 0; i < rdc_max_sets; i++) {
341         krdc = &rdc_k_info[i];
342         bzero(krdc, sizeof (*krdc));
343         krdc->index = i;
344         mutex_init(&krdc->dc_sleep, NULL, MUTEX_DRIVER, NULL);
345         mutex_init(&krdc->bmappmutex, NULL, MUTEX_DRIVER, NULL);
346         mutex_init(&krdc->kstat_mutex, NULL, MUTEX_DRIVER, NULL);
347         mutex_init(&krdc->bmp_kstat_mutex, NULL, MUTEX_DRIVER, NULL);
348         mutex_init(&krdc->syncbitmutex, NULL, MUTEX_DRIVER, NULL);
349         cv_init(&krdc->busycv, NULL, CV_DRIVER, NULL);
350         cv_init(&krdc->closingcv, NULL, CV_DRIVER, NULL);
351         cv_init(&krdc->haltcv, NULL, CV_DRIVER, NULL);
352         cv_init(&krdc->synccv, NULL, CV_DRIVER, NULL);
353     }

355     rdc_volume_update = nsc_register_svc("RDCVolumeUpdated",

```

```

356         rdc_volume_update_svc);
358     return (0);
359 }
_____unchanged_portion_omitted_____

```

```

*****
16833 Thu Oct 23 11:04:51 2014
new/usr/src/uts/common/avs/ns/solaris/nsc_raw.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

117 int
118 _nsc_init_raw(int maxdevs)
119 {
120     _nsc_raw_files =
121         kmem_zalloc(sizeof (*_nsc_raw_files) * maxdevs, KM_SLEEP);
122     if (!_nsc_raw_files)
123         return (ENOMEM);

123     _nsc_raw_maxdevs = maxdevs;
124     _nsc_raw_majors = NULL;

126     mutex_init(&_nsc_raw_lock, NULL, MUTEX_DRIVER, NULL);
127     return (0);
128 }
_____unchanged_portion_omitted_____

275 /*
276 * _raw_open
277 *
278 * Multiple opens, single close.
279 */

281 /* ARGSUSED */
282 static int
283 _raw_open(char *path, int flag, blind_t *cdp, void *iodev)
284 {
285     struct cred *cred;
286     raw_dev_t *cdi = NULL;
287     char *spath;
288     dev_t rdev;
289     int rc, cd, the_cd;
290     int plen;
291     ldi_ident_t li;

293     if (proc_nskernd == NULL) {
294         cmn_err(CE_WARN, "nskern: no nskernd daemon running!");
295         return (ENXIO);
296     }

298     if (_nsc_raw_maxdevs == 0) {
299         cmn_err(CE_WARN, "nskern: _raw_open() before _nsc_init_raw(!)");
300         return (ENXIO);
301     }

303     plen = strlen(path) + 1;
304     spath = kmem_alloc(plen, KM_SLEEP);
305     if (spath == NULL) {
306         cmn_err(CE_WARN,
307             "nskern: unable to alloc memory in _raw_open()");
308         return (ENOMEM);
309     }

306     (void) strcpy(spath, path);

308     /*
309     * Lookup the vnode to extract the dev_t info,
310     * then release the vnode.

```

```

311     /*
312     if ((rdev = ldi_get_dev_t_from_path(path)) == 0) {
313         kmem_free(spath, plen);
314         return (ENXIO);
315     }

317     /*
318     * See if this device is already opened
319     */

321     the_cd = -1;

323     mutex_enter(&_nsc_raw_lock);

325     for (cd = 0, cdi = _nsc_raw_files; cd < fd_hwm; cd++, cdi++) {
326         if (rdev == cdi->rdev) {
327             the_cd = cd;
328             break;
329         } else if (the_cd == -1 && !cdi->in_use)
330             the_cd = cd;
331     }

333     if (the_cd == -1) {
334         if (fd_hwm < _nsc_raw_maxdevs)
335             the_cd = fd_hwm++;
336         else {
337             mutex_exit(&_nsc_raw_lock);
338             cmn_err(CE_WARN, "_raw_open: too many open devices");
339             kmem_free(spath, plen);
340             return (EIO);
341         }
342     }

344     cdi = &_nsc_raw_files[the_cd];
345     if (cdi->in_use) {
346         /* already set up - just return */
347         mutex_exit(&_nsc_raw_lock);
348         *cdp = (blind_t)cdi->rdev;
349         kmem_free(spath, plen);
350         return (0);
351     }

353     cdi->partition = -1;
354     cdi->size = (uint64_t)0;
355     cdi->rdev = rdev;
356     cdi->path = spath;
357     cdi->plen = plen;

359     cred = ddi_get_cred();

361     /*
362     * Layered driver
363     *
364     * We use xxx_open_by_dev() since this guarantees that a
365     * specfs vnode is created and used, not a standard filesystem
366     * vnode. This is necessary since in a cluster PXFS will block
367     * vnode operations during switchovers, so we have to use the
368     * underlying specfs vnode not the PXFS vnode.
369     */

372     if ((rc = ldi_ident_from_dev(cdi->rdev, &li)) == 0) {
373         rc = ldi_open_by_dev(&cdi->rdev,
374             OTYP_BLK, FREAD|FWRITE, cred, &cdi->lh, li);
375     }
376     if (rc != 0) {

```

```
377         cdi->lh = NULL;
378         goto failed;
379     }
381     /*
382     * grab the major_t related information
383     */
385     cdi->major = _raw_get_maj_info(getmajor(rdev));
386     if (cdi->major == NULL) {
387         /* Out of memory */
388         cmn_err(CE_WARN,
389              "_raw_open: cannot alloc major number structure");
391         rc = ENOMEM;
392         goto failed;
393     }
395     *cdp = (blind_t)cdi->rdev;
396     cdi->in_use++;
398     mutex_exit(&nsc_raw_lock);
400     return (rc);
402 failed:
404     if (cdi->lh)
405         (void) ldi_close(cdi->lh, FWRITE|FREAD, cred);
407     bzero(cdi, sizeof (*cdi));
409     mutex_exit(&nsc_raw_lock);
411     kmem_free(spath, plen);
412     return (rc);
413 }
_____unchanged_portion_omitted_____
```

```

*****
60696 Thu Oct 23 11:04:51 2014
new/usr/src/uts/common/avs/ns/sv/sv.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

```

```

1727 static int
1728 sviocctl(dev_t dev, int cmd, intptr_t arg, int mode, cred_t *crp, int *rvalp)
1729 {
1730     char itmp1[12], itmp2[12]; /* temp char array for editing ints */
1731     spcs_s_info_t kstatus; /* Kernel version of spcs status */
1732     spcs_s_info_t ustatus; /* Address of user version of spcs status */
1733     sv_list32_t svl32; /* 32 bit Initial structure for SVIOC_LIST */
1734     sv_version_t svv; /* Version structure */
1735     sv_conf_t svc; /* User config structure */
1736     sv_list_t svl; /* Initial structure for SVIOC_LIST */
1737     void *usvn; /* Address of user sv_name_t */
1738     void *svn = NULL; /* Array for SVIOC_LIST */
1739     uint64_t phash; /* pathname hash */
1740     int rc = 0; /* Return code -- errno */
1741     int size; /* Number of items in array */
1742     int bytes; /* Byte size of array */
1743     int ilp32; /* Convert data structures for ilp32 userland */

1745     *rvalp = 0;

1747     /*
1748     * If sv_mod_status is 0 or SV_PREVENT_UNLOAD, then it will continue.
1749     * else it means it previously was SV_PREVENT_UNLOAD, and now it's
1750     * SV_ALLOW_UNLOAD, expecting the driver to eventually unload.
1751     *
1752     * SV_ALLOW_UNLOAD is final state, so no need to grab sv_mutex.
1753     */
1754     if (sv_mod_status == SV_ALLOW_UNLOAD) {
1755         return (EBUSY);
1756     }

1758     if ((cmd != SVIOC_LIST) && ((rc = drv_priv(crp)) != 0))
1759         return (rc);

1761     kstatus = spcs_s_kcreate();
1762     if (!kstatus) {
1763         DTRACE_PROBE1(sv_ioctl_err_kcreate, dev_t, dev);
1764         return (ENOMEM);
1765     }

1767     ilp32 = (ddi_model_convert_from((mode & FMODELS)) == DDI_MODEL_ILP32);

1769     switch (cmd) {
1771     case SVIOC_ENABLE:
1773         if (ilp32) {
1774             sv_conf32_t svc32;
1776             if (ddi_copyin((void *)arg, &svc32,
1777                 sizeof (svc32), mode) < 0) {
1778                 spcs_s_kfree(kstatus);
1779                 return (EFAULT);
1780             }

1782             svc.svc_error = (spcs_s_info_t)svc32.svc_error;
1783             (void) strcpy(svc.svc_path, svc32.svc_path);

```

```

1784             svc.svc_flag = svc32.svc_flag;
1785             svc.svc_major = svc32.svc_major;
1786             svc.svc_minor = svc32.svc_minor;
1787         } else {
1788             if (ddi_copyin((void *)arg, &svc,
1789                 sizeof (svc), mode) < 0) {
1790                 spcs_s_kfree(kstatus);
1791                 return (EFAULT);
1792             }
1793         }

1795     /* force to raw access */
1796     svc.svc_flag = NSC_DEVICE;

1798     if (sv_tset == NULL) {
1799         mutex_enter(&sv_mutex);
1801         if (sv_tset == NULL) {
1802             sv_tset = nst_init("sv_thr", sv_threads);
1803         }
1805         mutex_exit(&sv_mutex);

1807         if (sv_tset == NULL) {
1808             cmn_err(CE_WARN,
1809                 "!sv: could not allocate %d threads",
1810                 sv_threads);
1811         }
1812     }

1814     rc = sv_enable(svc.svc_path, svc.svc_flag,
1815         makedevice(svc.svc_major, svc.svc_minor), kstatus);

1817     if (rc == 0) {
1818         sv_config_time = nsc_lbolt();

1820         mutex_enter(&sv_mutex);
1821         sv_thread_tune(sv_threads_dev);
1822         mutex_exit(&sv_mutex);
1823     }

1825     DTRACE_PROBE3(sv_ioctl_end, dev_t, dev, int, *rvalp, int, rc);

1827     return (spcs_s_ocopyoutf(&kstatus, svc.svc_error, rc));
1828     /* NOTREACHED */

1830     case SVIOC_DISABLE:
1832         if (ilp32) {
1833             sv_conf32_t svc32;
1835             if (ddi_copyin((void *)arg, &svc32,
1836                 sizeof (svc32), mode) < 0) {
1837                 spcs_s_kfree(kstatus);
1838                 return (EFAULT);
1839             }

1841             svc.svc_error = (spcs_s_info_t)svc32.svc_error;
1842             svc.svc_major = svc32.svc_major;
1843             svc.svc_minor = svc32.svc_minor;
1844             (void) strcpy(svc.svc_path, svc32.svc_path);
1845             svc.svc_flag = svc32.svc_flag;
1846         } else {
1847             if (ddi_copyin((void *)arg, &svc,
1848                 sizeof (svc), mode) < 0) {
1849                 spcs_s_kfree(kstatus);

```

```

1850         return (EFAULT);
1851     }
1852 }
1854 if (svc.svc_major == (major_t)-1 &&
1855     svc.svc_minor == (minor_t)-1) {
1856     sv_dev_t *svp;
1857     int i;
1859     /*
1860      * User level could not find the minor device
1861      * node, so do this the slow way by searching
1862      * the entire sv config for a matching pathname.
1863      */
1865     phash = nsc_strhash(svc.svc_path);
1867     mutex_enter(&sv_mutex);
1869     for (i = 0; i < sv_max_devices; i++) {
1870         svp = &sv_devs[i];
1872         if (svp->sv_state == SV_DISABLE ||
1873             svp->sv_fd == NULL)
1874             continue;
1876         if (nsc_fdpathcmp(svp->sv_fd, phash,
1877             svc.svc_path) == 0) {
1878             svc.svc_major = getmajor(svp->sv_dev);
1879             svc.svc_minor = getminor(svp->sv_dev);
1880             break;
1881         }
1882     }
1884     mutex_exit(&sv_mutex);
1886     if (svc.svc_major == (major_t)-1 &&
1887         svc.svc_minor == (minor_t)-1)
1888         return (spcs_s_ocopyoutf(&kstatus,
1889             svc.svc_error, SV_ENODEV));
1890 }
1892 rc = sv_disable(makedevice(svc.svc_major, svc.svc_minor),
1893     kstatus);
1895 if (rc == 0) {
1896     sv_config_time = nsc_lbolt();
1898     mutex_enter(&sv_mutex);
1899     sv_thread_tune(-sv_threads_dev);
1900     mutex_exit(&sv_mutex);
1901 }
1903 DTRACE_PROBE3(sv_ioctl_2, dev_t, dev, int, *rvalp, int, rc);
1905 return (spcs_s_ocopyoutf(&kstatus, svc.svc_error, rc));
1906 /* NOTREACHED */
1908 case SVIOC_LIST:
1910     if (ilp32) {
1911         if (ddi_copyin((void *)arg, &svl32,
1912             sizeof (svl32), mode) < 0) {
1913             spcs_s_kfree(kstatus);
1914             return (EFAULT);
1915         }

```

```

1917         ustatus = (spcs_s_info_t)svl32.svl_error;
1918         size = svl32.svl_count;
1919         usvn = (void *) (unsigned long)svl32.svl_names;
1920     } else {
1921         if (ddi_copyin((void *)arg, &svl,
1922             sizeof (svl), mode) < 0) {
1923             spcs_s_kfree(kstatus);
1924             return (EFAULT);
1925         }
1927         ustatus = svl.svl_error;
1928         size = svl.svl_count;
1929         usvn = svl.svl_names;
1930     }
1932     /* Do some boundary checking */
1933     if ((size < 0) || (size > sv_max_devices)) {
1934         /* Array size is out of range */
1935         return (spcs_s_ocopyoutf(&kstatus, ustatus,
1936             SV_EARBOUNDS, "0",
1937             spcs_s_inttostring(sv_max_devices, itmpl,
1938                 sizeof (itmpl), 0),
1939             spcs_s_inttostring(size, itmp2,
1940                 sizeof (itmp2), 0)));
1941     }
1943     if (ilp32)
1944         bytes = size * sizeof (sv_name32_t);
1945     else
1946         bytes = size * sizeof (sv_name_t);
1948     /* Allocate memory for the array of structures */
1949     if (bytes != 0) {
1950         svn = kmem_zalloc(bytes, KM_SLEEP);
1951         if (!svn) {
1952             return (spcs_s_ocopyoutf(&kstatus,
1953                 ustatus, ENOMEM));
1954         }
1955     }
1956     rc = sv_list(svn, size, rvalp, ilp32);
1957     if (rc) {
1958         if (svn != NULL)
1959             kmem_free(svn, bytes);
1960         return (spcs_s_ocopyoutf(&kstatus, ustatus, rc));
1961     }
1962     if (ilp32) {
1963         svl32.svl_timestamp = (uint32_t)sv_config_time;
1964         svl32.svl_maxdevs = (int32_t)sv_max_devices;
1965     }
1966     /* Return the list structure */
1967     if (ddi_copyout(&svl32, (void *)arg,
1968         sizeof (svl32), mode) < 0) {
1969         spcs_s_kfree(kstatus);
1970         if (svn != NULL)
1971             kmem_free(svn, bytes);
1972         return (EFAULT);
1973     }
1974 } else {
1975     svl.svl_timestamp = sv_config_time;
1976     svl.svl_maxdevs = sv_max_devices;
1977     /* Return the list structure */
1978     if (ddi_copyout(&svl, (void *)arg,

```

```

1978         sizeof (svl), mode) < 0) {
1979             spcs_s_kfree(kstatus);
1980             if (svn != NULL)
1981                 kmem_free(svn, bytes);
1982             return (EFAULT);
1983         }
1984     }
1985
1986     /* Return the array */
1987     if (svn != NULL) {
1988         if (ddi_copyout(svn, usvn, bytes, mode) < 0) {
1989             kmem_free(svn, bytes);
1990             spcs_s_kfree(kstatus);
1991             return (EFAULT);
1992         }
1993         kmem_free(svn, bytes);
1994     }
1995
1996     DTRACE_PROBE3(sv_ioctl_3, dev_t, dev, int, *rvalp, int, 0);
1997
1998     return (spcs_s_ocopyoutf(&kstatus, ustatus, 0));
1999     /* NOTREACHED */
2000
2001     case SVIOC_VERSION:
2002
2003         if (ilp32) {
2004             sv_version32_t svv32;
2005
2006             if (ddi_copyin((void *)arg, &svv32,
2007                 sizeof (svv32), mode) < 0) {
2008                 spcs_s_kfree(kstatus);
2009                 return (EFAULT);
2010             }
2011
2012             svv32.svv_major_rev = sv_major_rev;
2013             svv32.svv_minor_rev = sv_minor_rev;
2014             svv32.svv_micro_rev = sv_micro_rev;
2015             svv32.svv_baseline_rev = sv_baseline_rev;
2016
2017             if (ddi_copyout(&svv32, (void *)arg,
2018                 sizeof (svv32), mode) < 0) {
2019                 spcs_s_kfree(kstatus);
2020                 return (EFAULT);
2021             }
2022
2023             ustatus = (spcs_s_info_t)svv32.svv_error;
2024         } else {
2025             if (ddi_copyin((void *)arg, &svv,
2026                 sizeof (svv), mode) < 0) {
2027                 spcs_s_kfree(kstatus);
2028                 return (EFAULT);
2029             }
2030
2031             svv.svv_major_rev = sv_major_rev;
2032             svv.svv_minor_rev = sv_minor_rev;
2033             svv.svv_micro_rev = sv_micro_rev;
2034             svv.svv_baseline_rev = sv_baseline_rev;
2035
2036             if (ddi_copyout(&svv, (void *)arg,
2037                 sizeof (svv), mode) < 0) {
2038                 spcs_s_kfree(kstatus);
2039                 return (EFAULT);
2040             }
2041
2042             ustatus = svv.svv_error;
2043         }

```

```

2045         DTRACE_PROBE3(sv_ioctl_4, dev_t, dev, int, *rvalp, int, 0);
2046
2047         return (spcs_s_ocopyoutf(&kstatus, ustatus, 0));
2048         /* NOTREACHED */
2049
2050     case SVIOC_UNLOAD:
2051         rc = sv_prepare_unload();
2052
2053         if (ddi_copyout(&rc, (void *)arg, sizeof (rc), mode) < 0) {
2054             rc = EFAULT;
2055         }
2056
2057         spcs_s_kfree(kstatus);
2058         return (rc);
2059
2060     default:
2061         spcs_s_kfree(kstatus);
2062
2063         DTRACE_PROBE3(sv_ioctl_4, dev_t, dev, int, *rvalp, int, EINVAL);
2064
2065         return (EINVAL);
2066         /* NOTREACHED */
2067     }
2068
2069     /* NOTREACHED */
2070 }

```

unchanged_portion_omitted

```

*****
128510 Thu Oct 23 11:04:52 2014
new/usr/src/uts/common/crypto/io/dca.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

```

```

1587 dca_request_t *
1588 dca_newreq(dca_t *dca)
1589 {
1590     dca_request_t      *reqp;
1591     size_t              size;
1592     ddi_dma_cookie_t    c;
1593     unsigned            nc;
1594     int                 rv;
1595     int                 n_chain = 0;
1597     size = (DESC_SIZE * MAXFRAGS) + CTX_MAXLENGTH;
1599     reqp = kmem_zalloc(sizeof (dca_request_t), KM_SLEEP);
1601     reqp->dr_dca = dca;
1603     /*
1604      * Setup the DMA region for the context and descriptors.
1605      */
1606     rv = ddi_dma_alloc_handle(dca->dca_dip, &dca_dmaattr, DDI_DMA_SLEEP,
1607         NULL, &reqp->dr_ctx_dmah);
1608     if (rv != DDI_SUCCESS) {
1609         dca_error(dca, "failure allocating request DMA handle");
1610         dca_destroyreq(reqp);
1611         return (NULL);
1612     }
1614     /* for driver hardening, allocate in whole pages */
1615     rv = ddi_dma_mem_alloc(reqp->dr_ctx_dmah,
1616         ROUNDUP(size, dca->dca_pagesize), &dca_devattr, DDI_DMA_CONSISTENT,
1617         DDI_DMA_SLEEP, NULL, &reqp->dr_ctx_kaddr, &size,
1618         &reqp->dr_ctx_acch);
1619     if (rv != DDI_SUCCESS) {
1620         dca_error(dca, "unable to alloc request DMA memory");
1621         dca_destroyreq(reqp);
1622         return (NULL);
1623     }
1625     rv = ddi_dma_addr_bind_handle(reqp->dr_ctx_dmah, NULL,
1626         reqp->dr_ctx_kaddr, size, DDI_DMA_CONSISTENT | DDI_DMA_WRITE,
1627         DDI_DMA_SLEEP, 0, &c, &nc);
1628     if (rv != DDI_DMA_MAPPED) {
1629         dca_error(dca, "failed binding request DMA handle");
1630         dca_destroyreq(reqp);
1631         return (NULL);
1632     }
1633     reqp->dr_ctx_paddr = c.dmac_address;
1635     reqp->dr_dma_size = size;
1637     /*
1638      * Set up the dma for our scratch/shared buffers.
1639      */
1640     rv = ddi_dma_alloc_handle(dca->dca_dip, &dca_dmaattr,
1641         DDI_DMA_SLEEP, NULL, &reqp->dr_ibuf_dmah);
1642     if (rv != DDI_SUCCESS) {
1643         dca_error(dca, "failure allocating ibuf DMA handle");
1644         dca_destroyreq(reqp);

```

```

1645         return (NULL);
1646     }
1647     rv = ddi_dma_alloc_handle(dca->dca_dip, &dca_dmaattr,
1648         DDI_DMA_SLEEP, NULL, &reqp->dr_obuf_dmah);
1649     if (rv != DDI_SUCCESS) {
1650         dca_error(dca, "failure allocating obuf DMA handle");
1651         dca_destroyreq(reqp);
1652         return (NULL);
1653     }
1655     rv = ddi_dma_alloc_handle(dca->dca_dip, &dca_dmaattr,
1656         DDI_DMA_SLEEP, NULL, &reqp->dr_chain_in_dmah);
1657     if (rv != DDI_SUCCESS) {
1658         dca_error(dca, "failure allocating chain_in DMA handle");
1659         dca_destroyreq(reqp);
1660         return (NULL);
1661     }
1663     rv = ddi_dma_alloc_handle(dca->dca_dip, &dca_dmaattr,
1664         DDI_DMA_SLEEP, NULL, &reqp->dr_chain_out_dmah);
1665     if (rv != DDI_SUCCESS) {
1666         dca_error(dca, "failure allocating chain_out DMA handle");
1667         dca_destroyreq(reqp);
1668         return (NULL);
1669     }
1671     /*
1672      * for driver hardening, allocate in whole pages.
1673      */
1674     size = ROUNDUP(MAXPACKET, dca->dca_pagesize);
1675     #if defined(i386) || defined(__i386)
1676     /*
1677      * Use kmem_alloc instead of ddi_dma_mem_alloc here since the latter
1678      * may fail on x86 platform if a physically contiguous memory chunk
1679      * cannot be found. From initial testing, we did not see performance
1680      * degradation as seen on Sparc.
1681      */
1682     reqp->dr_ibuf_kaddr = kmem_alloc(size, KM_SLEEP);
1683     reqp->dr_obuf_kaddr = kmem_alloc(size, KM_SLEEP);
1684     if ((reqp->dr_ibuf_kaddr = kmem_alloc(size, KM_SLEEP)) == NULL) {
1685         dca_error(dca, "unable to alloc request ibuf memory");
1686         dca_destroyreq(reqp);
1687         return (NULL);
1688     }
1689     if ((reqp->dr_obuf_kaddr = kmem_alloc(size, KM_SLEEP)) == NULL) {
1690         dca_error(dca, "unable to alloc request obuf memory");
1691         dca_destroyreq(reqp);
1692         return (NULL);
1693     }
1694     #else
1695     /*
1696      * We could kmem_alloc for Sparc too. However, it gives worse
1697      * performance when transferring more than one page data. For example,
1698      * using 4 threads and 12032 byte data and 3DES on 900MHZ Sparc system,
1699      * kmem_alloc uses 80% CPU and ddi_dma_mem_alloc uses 50% CPU for
1700      * the same throughput.
1701     */

```



```
1702     rv = ddi_dma_mem_alloc(reqp->dr_obuf_dmah,
1703     size, &dca_bufattr,
1704     DDI_DMA_STREAMING, DDI_DMA_SLEEP, NULL, &reqp->dr_obuf_kaddr,
1705     &size, &reqp->dr_obuf_acch);
1706     if (rv != DDI_SUCCESS) {
1707         dca_error(dca, "unable to alloc request DMA memory");
1708         dca_destroyreq(reqp);
1709         return (NULL);
1710     }
1711 #endif

1713     /* Skip the used portion in the context page */
1714     reqp->dr_offset = CTX_MAXLENGTH;
1715     if ((rv = dca_bindchains_one(reqp, size, reqp->dr_offset,
1716     reqp->dr_ibuf_kaddr, reqp->dr_ibuf_dmah,
1717     DDI_DMA_WRITE | DDI_DMA_STREAMING,
1718     &reqp->dr_ibuf_head, &n_chain)) != DDI_SUCCESS) {
1719         (void) dca_destroyreq(reqp);
1720         return (NULL);
1721     }
1722     reqp->dr_ibuf_paddr = reqp->dr_ibuf_head.dc_buffer_paddr;
1723     /* Skip the space used by the input buffer */
1724     reqp->dr_offset += DESC_SIZE * n_chain;

1726     if ((rv = dca_bindchains_one(reqp, size, reqp->dr_offset,
1727     reqp->dr_obuf_kaddr, reqp->dr_obuf_dmah,
1728     DDI_DMA_READ | DDI_DMA_STREAMING,
1729     &reqp->dr_obuf_head, &n_chain)) != DDI_SUCCESS) {
1730         (void) dca_destroyreq(reqp);
1731         return (NULL);
1732     }
1733     reqp->dr_obuf_paddr = reqp->dr_obuf_head.dc_buffer_paddr;
1734     /* Skip the space used by the output buffer */
1735     reqp->dr_offset += DESC_SIZE * n_chain;

1737     DBG(dca, DCHATTY, "CTX is 0x%p, phys 0x%x, len %d",
1738     reqp->dr_ctx_kaddr, reqp->dr_ctx_paddr, CTX_MAXLENGTH);
1739     return (reqp);
1740 }
unchanged_portion_omitted
```

```

*****
7815 Thu Oct 23 11:04:52 2014
new/usr/src/uts/common/crypto/io/dca_rng.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */

22 /*
23  * Copyright 2005 Sun Microsystems, Inc. All rights reserved.
24  * Use is subject to license terms.
25 */

27 #pragma ident      "%Z%M% %I%      %E% SMI"

27 /*
28  * Deimos - cryptographic acceleration based upon Broadcom 582x.
29  */

31 #include <sys/types.h>
32 #include <sys/ddi.h>
33 #include <sys/sunddi.h>
34 #include <sys/kmem.h>
35 #include <sys/crypto/dca.h>
36 #include <sys/atomic.h>

38 /*
39  * Random number implementation.
40  */

42 static int dca_rngstart(dca_t *, dca_request_t *);
43 static void dca_rngdone(dca_request_t *, int);

45 static void dca_random_done();
46 int dca_random_buffer(dca_t *dca, caddr_t buf, int len);
47 int dca_random_init();
48 void dca_random_fini();

50 int
51 dca_rng(dca_t *dca, uchar_t *buf, size_t len, crypto_req_handle_t req)
52 {
53     dca_request_t     *reqp;
54     int                rv;
55     crypto_data_t     *data;

57     if ((reqp = dca_getreq(dca, MCR2, 1)) == NULL) {
58         dca_error(dca, "unable to allocate request for RNG");

```

```

59         return (CRYPTO_HOST_MEMORY);
60     }

62     reqp->dr_kcf_req = req;

64     data = &reqp->dr_ctx.in_dup;
65     data->cd_format = CRYPTO_DATA_RAW;
66     data->cd_offset = 0;
67     data->cd_length = 0;
68     data->cd_raw.iov_base = (char *)buf;
69     data->cd_raw.iov_len = len;
70     reqp->dr_out = data;
71     reqp->dr_in = NULL;

73     rv = dca_rngstart(dca, reqp);
74     if (rv != CRYPTO_QUEUED) {
75         if (reqp->destroy)
76             dca_destroyreq(reqp);
77         else
78             dca_freereq(reqp);
79     }
80     return (rv);
81 }

unchanged_portion_omitted_

202 /*
203  * This gives a 32k random bytes per buffer. The two buffers will switch back
204  * and forth. When a buffer is used up, a request will be submitted to refill
205  * this buffer before switching to the other one
206  */

208 #define RANDOM_BUFFER_SIZE          (1<<15)
209 #define DCA_RANDOM_MAX_WAIT         10000

211 int
212 dca_random_init(dca_t *dca)
213 {
214     /* Mutex for the local random number pool */
215     mutex_init(&dca->dca_random_lock, NULL, MUTEX_DRIVER, NULL);

217     dca->dca_buf1 = kmem_alloc(RANDOM_BUFFER_SIZE, KM_SLEEP);
219     if ((dca->dca_buf1 = kmem_alloc(RANDOM_BUFFER_SIZE, KM_SLEEP)) ==
220         NULL) {
221         mutex_destroy(&dca->dca_random_lock);
222         return (CRYPTO_FAILED);
223     }

219     dca->dca_buf2 = kmem_alloc(RANDOM_BUFFER_SIZE, KM_SLEEP);
225     if ((dca->dca_buf2 = kmem_alloc(RANDOM_BUFFER_SIZE, KM_SLEEP)) ==
226         NULL) {
227         mutex_destroy(&dca->dca_random_lock);
228         kmem_free(dca->dca_buf1, RANDOM_BUFFER_SIZE);
229         return (CRYPTO_FAILED);
230     }

221     return (CRYPTO_SUCCESS);
222 }

unchanged_portion_omitted_

```

```

*****
298742 Thu Oct 23 11:04:52 2014
new/usr/src/uts/common/crypto/io/dprov.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

9102 /*
9103  * Create an object from the specified template. Checks whether the
9104  * object can be created according to its attributes and the state
9105  * of the session. The new session object id is returned. If the
9106  * object is a token object, it is added to the per-instance object
9107  * table as well.
9108  */
9109 static int
9110 dprov_create_object_from_template(dprov_state_t *softc,
9111     dprov_session_t *session, crypto_object_attribute_t *template,
9112     uint_t nattr, crypto_object_id_t *object_id, boolean_t check_for_secret,
9113     boolean_t force)
9114 {
9115     dprov_object_t *object;
9116     boolean_t is_token = B_FALSE;
9117     boolean_t extractable_attribute_present = B_FALSE;
9118     boolean_t sensitive_attribute_present = B_FALSE;
9119     boolean_t private_attribute_present = B_FALSE;
9120     boolean_t token_attribute_present = B_FALSE;
9121     uint_t i;
9122     int error;
9123     uint_t attr;
9124     uint_t oattr;
9125     crypto_attr_type_t type;
9126     size_t old_len, new_len;
9127     offset_t offset;

9129     if (nattr > DPROV_MAX_ATTR)
9130         return (CRYPTO_HOST_MEMORY);

9132     if (!force) {
9133         /* verify that object can be created */
9134         if ((error = dprov_template_can_create(session, template,
9135             nattr, check_for_secret)) != CRYPTO_SUCCESS)
9136             return (error);
9137     }

9139     /* allocate new object */
9140     object = kmem_zalloc(sizeof (dprov_object_t), KM_SLEEP);
9141     if (object == NULL)
9142         return (CRYPTO_HOST_MEMORY);

9142     /* is it a token object? */
9143     /* check CKA_TOKEN attribute value */
9144     error = dprov_get_template_attr_boolean(template, nattr,
9145         DPROV_CKA_TOKEN, &is_token);
9146     if (error == CRYPTO_SUCCESS && is_token) {
9147         /* token object, add it to the per-instance object table */
9148         for (i = 0; i < DPROV_MAX_OBJECTS; i++)
9149             if (softc->ds_objects[i] == NULL)
9150                 break;
9151         if (i == DPROV_MAX_OBJECTS)
9152             /* no free slot */
9153             return (CRYPTO_HOST_MEMORY);
9154         softc->ds_objects[i] = object;
9155         object->do_token_idx = i;
9156         DPROV_OBJECT_REFHOLD(object);
9157     }

```

```

9159     /* add object to session object table */
9160     for (i = 0; i < DPROV_MAX_OBJECTS; i++)
9161         if (session->ds_objects[i] == NULL)
9162             break;
9163     if (i == DPROV_MAX_OBJECTS) {
9164         /* no more session object slots */
9165         DPROV_OBJECT_REFRELE(object);
9166         return (CRYPTO_HOST_MEMORY);
9167     }
9168     session->ds_objects[i] = object;
9169     DPROV_OBJECT_REFHOLD(object);
9170     *object_id = i;

9172     /* initialize object from template */
9173     for (attr = 0, oattr = 0; attr < nattr; attr++) {
9174         if (template[attr].oa_value == NULL)
9175             continue;
9176         type = template[attr].oa_type;
9177         old_len = template[attr].oa_value_len;
9178         new_len = attribute_size(type, old_len);

9180         if (type == DPROV_CKA_EXTRACTABLE) {
9181             extractable_attribute_present = B_TRUE;
9182         } else if (type == DPROV_CKA_PRIVATE) {
9183             private_attribute_present = B_TRUE;
9184         } else if (type == DPROV_CKA_TOKEN) {
9185             token_attribute_present = B_TRUE;
9186         }
9187         object->do_attr[oattr].oa_type = type;
9188         object->do_attr[oattr].oa_value_len = new_len;

9190         object->do_attr[oattr].oa_value = kmem_zalloc(new_len,
9191             KM_SLEEP);

9193         offset = 0;
9194 #ifdef _BIG_ENDIAN
9195         if (fixed_size_attribute(type)) {
9196             offset = old_len - new_len;
9197         }
9198 #endif
9199         bcopy(&template[attr].oa_value[offset],
9200             object->do_attr[oattr].oa_value, new_len);
9201         oattr++;
9202     }

9204     /* add boolean attributes that must be present */
9205     if (extractable_attribute_present == B_FALSE) {
9206         object->do_attr[oattr].oa_type = DPROV_CKA_EXTRACTABLE;
9207         object->do_attr[oattr].oa_value_len = 1;
9208         object->do_attr[oattr].oa_value = kmem_alloc(1, KM_SLEEP);
9209         object->do_attr[oattr].oa_value[0] = B_TRUE;
9210         oattr++;
9211     }

9213     if (private_attribute_present == B_FALSE) {
9214         object->do_attr[oattr].oa_type = DPROV_CKA_PRIVATE;
9215         object->do_attr[oattr].oa_value_len = 1;
9216         object->do_attr[oattr].oa_value = kmem_alloc(1, KM_SLEEP);
9217         object->do_attr[oattr].oa_value[0] = B_FALSE;
9218         oattr++;
9219     }

9221     if (token_attribute_present == B_FALSE) {
9222         object->do_attr[oattr].oa_type = DPROV_CKA_TOKEN;
9223         object->do_attr[oattr].oa_value_len = 1;

```

```
9224         object->do_attr[oattr].oa_value = kmem_alloc(1, KM_SLEEP);
9225         object->do_attr[oattr].oa_value[0] = B_FALSE;
9226         oattr++;
9227     }

9229     if (sensitive_attribute_present == B_FALSE) {
9230         object->do_attr[oattr].oa_type = DPROV_CKA_SENSITIVE;
9231         object->do_attr[oattr].oa_value_len = 1;
9232         object->do_attr[oattr].oa_value = kmem_alloc(1, KM_SLEEP);
9233         object->do_attr[oattr].oa_value[0] = B_FALSE;
9234         oattr++;
9235     }
9236     return (CRYPTO_SUCCESS);
9237 }
_____unchanged_portion_omitted_____
```

```

*****
10050 Thu Oct 23 11:04:53 2014
new/usr/src/uts/common/fs/nfs/nfs_sys.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

92 int
93 nfssys(enum nfssys_op opcode, void *arg)
94 {
95     int error = 0;

97     if (!(opcode == NFS_REVAUTH || opcode == NFS4_SVC) &&
98         secpolicy_nfs(CRED()) != 0)
99         return (set_errno(EPERM));

101     switch (opcode) {
102     case NFS4_CLR_STATE: { /* Clear NFS4 client state */
103         struct nfs4clrst_args clr;
104         STRUCT_DECL(nfs4clrst_args, u_clr);

106         /*
107          * If the server is not loaded then no point in
108          * clearing nothing :- )
109          */
110         if (rfs4_client_clrst == NULL) {
111             break;
112         }

114         if (!INGLOBALZONE(curproc))
115             return (set_errno(EPERM));

117         STRUCT_INIT(u_clr, get_udatamodel());

119         if (copyin(arg, STRUCT_BUF(u_clr), STRUCT_SIZE(u_clr)))
120             return (set_errno(EFAULT));

122         clr.vers = STRUCT_FGET(u_clr, vers);

124         if (clr.vers != NFS4_CLRST_VERSION)
125             return (set_errno(EINVAL));

127         clr.addr_type = STRUCT_FGET(u_clr, addr_type);
128         clr.ap = STRUCT_FGETP(u_clr, ap);
129         rfs4_client_clrst(&clr);
130         break;
131     }

133     case SVCPOOL_CREATE: { /* setup an RPC server thread pool */
134         struct svcpool_args p;

136         if (copyin(arg, &p, sizeof(p)))
137             return (set_errno(EFAULT));

139         error = svc_pool_create(&p);
140         break;
141     }

143     case SVCPOOL_WAIT: { /* wait in kernel for threads to be needed */
144         int id;

146         if (copyin(arg, &id, sizeof(id)))
147             return (set_errno(EFAULT));

149         error = svc_wait(id);

```

```

150         break;
151     }

153     case SVCPOOL_RUN: { /* give work to a runnable thread */
154         int id;

156         if (copyin(arg, &id, sizeof(id)))
157             return (set_errno(EFAULT));

159         error = svc_do_run(id);
160         break;
161     }

163     case RDMA_SVC_INIT: {
164         struct rdma_svc_args rsa;
165         char netstore[20] = "tcp";

167         if (!INGLOBALZONE(curproc))
168             return (set_errno(EPERM));
169         if (get_udatamodel() != DATAMODEL_NATIVE) {
170             STRUCT_DECL(rdma_svc_args, urrsa);

172             STRUCT_INIT(urrsa, get_udatamodel());
173             if (copyin(arg, STRUCT_BUF(urrsa), STRUCT_SIZE(urrsa)))
174                 return (set_errno(EFAULT));

176             rsa.poolid = STRUCT_FGET(urrsa, poolid);
177             rsa.nfs_versmin = STRUCT_FGET(urrsa, nfs_versmin);
178             rsa.nfs_versmax = STRUCT_FGET(urrsa, nfs_versmax);
179             rsa.delegation = STRUCT_FGET(urrsa, delegation);
180         } else {
181             if (copyin(arg, &rsa, sizeof(rsa)))
182                 return (set_errno(EFAULT));
183         }
184         rsa.netid = netstore;

186         error = rdma_start(&rsa);
187         break;
188     }

190     case NFS_SVC: { /* NFS server daemon */
191         STRUCT_DECL(nfs_svc_args, nsa);

193         if (!INGLOBALZONE(curproc))
194             return (set_errno(EPERM));
195         STRUCT_INIT(nsa, get_udatamodel());

197         if (copyin(arg, STRUCT_BUF(nsa), STRUCT_SIZE(nsa)))
198             return (set_errno(EFAULT));

200         error = nfs_svc(STRUCT_BUF(nsa), get_udatamodel());
201         break;
202     }

204     case EXPORTFS: { /* export a file system */
205         error = nfs_export(arg);
206         break;
207     }

209     case NFS_GETFH: { /* get a file handle */
210         STRUCT_DECL(nfs_getfh_args, nga);

212         if (!INGLOBALZONE(curproc))
213             return (set_errno(EPERM));
214         STRUCT_INIT(nga, get_udatamodel());
215         if (copyin(arg, STRUCT_BUF(nga), STRUCT_SIZE(nga)))

```

```

216         return (set_errno(EFAULT));
218         error = nfs_getfh(STRUCT_BUF(nga), get_udatamodel(), CRED());
219         break;
220     }
222     case NFS_REVAUTH: { /* revoke the cached credentials for the uid */
223         STRUCT_DECL(nfs_revauth_args, nra);
225         STRUCT_INIT(nra, get_udatamodel());
226         if (copyin(arg, STRUCT_BUF(nra), STRUCT_SIZE(nra)))
227             return (set_errno(EFAULT));
229         /* This call performs its own privilege checking */
230         error = sec_clnt_revoke(STRUCT_FGET(nra, authtype),
231             STRUCT_FGET(nra, uid), CRED(), NULL, get_udatamodel());
232         break;
233     }
235     case LM_SVC: { /* LM server daemon */
236         struct lm_svc_args lsa;
238         if (get_udatamodel() != DATAMODEL_NATIVE) {
239             STRUCT_DECL(lm_svc_args, ulsa);
241             STRUCT_INIT(ulsa, get_udatamodel());
242             if (copyin(arg, STRUCT_BUF(ulsa), STRUCT_SIZE(ulsa)))
243                 return (set_errno(EFAULT));
245             lsa.version = STRUCT_FGET(ulsa, version);
246             lsa.fd = STRUCT_FGET(ulsa, fd);
247             lsa.n_fmly = STRUCT_FGET(ulsa, n_fmly);
248             lsa.n_proto = STRUCT_FGET(ulsa, n_proto);
249             lsa.n_rdev = expldev(STRUCT_FGET(ulsa, n_rdev));
250             lsa.debug = STRUCT_FGET(ulsa, debug);
251             lsa.timeout = STRUCT_FGET(ulsa, timeout);
252             lsa.grace = STRUCT_FGET(ulsa, grace);
253             lsa.retransmittimeout = STRUCT_FGET(ulsa,
254                 retransmittimeout);
255         } else {
256             if (copyin(arg, &lsa, sizeof(lsa)))
257                 return (set_errno(EFAULT));
258         }
260         error = lm_svc(&lsa);
261         break;
262     }
264     case KILL_LOCKMGR: {
265         error = lm_shutdown();
266         break;
267     }
269     case LOG_FLUSH: { /* Flush log buffer and possibly rename */
270         STRUCT_DECL(nfsl_flush_args, nfa);
272         STRUCT_INIT(nfa, get_udatamodel());
273         if (copyin(arg, STRUCT_BUF(nfa), STRUCT_SIZE(nfa)))
274             return (set_errno(EFAULT));
276         error = nfsl_flush(STRUCT_BUF(nfa), get_udatamodel());
277         break;
278     }
280     case NFS4_SVC: { /* NFS client callback daemon */

```

```

282         STRUCT_DECL(nfs4_svc_args, nsa);
284         STRUCT_INIT(nsa, get_udatamodel());
286         if (copyin(arg, STRUCT_BUF(nsa), STRUCT_SIZE(nsa)))
287             return (set_errno(EFAULT));
289         error = nfs4_svc(STRUCT_BUF(nsa), get_udatamodel());
290         break;
291     }
293     /* Request that NFSv4 server quiesce on next shutdown */
294     case NFS4_SVC_REQUEST_QUIESCE: {
295         int id;
297         /* check that nfssrv module is loaded */
298         if (nfs_srv_quiesce_func == NULL)
299             return (set_errno(ENOTSUP));
301         if (copyin(arg, &id, sizeof(id)))
302             return (set_errno(EFAULT));
304         error = svc_pool_control(id, SVCASET_SHUTDOWN_PROC,
305             (void *)nfs_srv_quiesce_func);
306         break;
307     }
309     case NFS_IDMAP: {
310         struct nfssidmap_args idm;
312         if (copyin(arg, &idm, sizeof(idm)))
313             return (set_errno(EFAULT));
315         nfs_idmap_args(&idm);
316         error = 0;
317         break;
318     }
320     case NFS4_DSS_SETPATHS_SIZE: {
321         /* crosses ILP32/LP64 boundary */
322         uint32_t nfs4_dss_bufsize = 0;
324         if (copyin(arg, &nfs4_dss_bufsize, sizeof(nfs4_dss_bufsize)))
325             return (set_errno(EFAULT));
326         nfs4_dss_bufflen = (long)nfs4_dss_bufsize;
327         error = 0;
328         break;
329     }
331     case NFS4_DSS_SETPATHS: {
332         char *nfs4_dss_bufp;
334         /* check that nfssrv module is loaded */
335         if (nfs_srv_dss_func == NULL)
336             return (set_errno(ENOTSUP));
338         /*
339          * NFS4_DSS_SETPATHS_SIZE must be called before
340          * NFS4_DSS_SETPATHS, to tell us how big a buffer we need
341          * to allocate.
342          */
343         if (nfs4_dss_bufflen == 0)
344             return (set_errno(EINVAL));
345         nfs4_dss_bufp = kmem_alloc(nfs4_dss_bufflen, KM_SLEEP);
346         if (nfs4_dss_bufp == NULL)
347             return (set_errno(ENOMEM));

```

```
347         if (copyin(arg, nfs4_dss_bufp, nfs4_dss_buflen)) {
348             kmem_free(nfs4_dss_bufp, nfs4_dss_buflen);
349             return (set_errno(EFAULT));
350         }
351
352         /* unpack the buffer and extract the pathnames */
353         error = nfs_srv_dss_func(nfs4_dss_bufp, nfs4_dss_buflen);
354         kmem_free(nfs4_dss_bufp, nfs4_dss_buflen);
355
356         break;
357     }
358
359     case NFS4_EPHEMERAL_MOUNT_TO: {
360         uint_t mount_to;
361
362         /*
363          * Not a very complicated call.
364          */
365         if (copyin(arg, &mount_to, sizeof (mount_to)))
366             return (set_errno(EFAULT));
367         nfs4_ephemeral_set_mount_to(mount_to);
368         error = 0;
369         break;
370     }
371
372     case MOUNTD_ARGS: {
373         uint_t did;
374
375         /*
376          * For now, only passing down the door fd; if we
377          * ever need to pass down more info, we can use
378          * a (properly aligned) struct.
379          */
380         if (copyin(arg, &did, sizeof (did)))
381             return (set_errno(EFAULT));
382         mountd_args(did);
383         error = 0;
384         break;
385     }
386
387     case NFSCMD_ARGS: {
388         uint_t did;
389
390         /*
391          * For now, only passing down the door fd; if we
392          * ever need to pass down more info, we can use
393          * a (properly aligned) struct.
394          */
395         if (copyin(arg, &did, sizeof (did)))
396             return (set_errno(EFAULT));
397         nfscmd_args(did);
398         error = 0;
399         break;
400     }
401
402     default:
403         error = EINVAL;
404         break;
405 }
406
407 return ((error != 0) ? set_errno(error) : 0);
408 }
```

unchanged portion omitted

new/usr/src/uts/common/fs/ufs/ufs_acl.c

1

```
*****
53529 Thu Oct 23 11:04:53 2014
new/usr/src/uts/common/fs/ufs/ufs_acl.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
unchanged_portion_omitted

712 /*ARGSUSED2*/
713 int
714 ufs_acl_get(struct inode *ip, vsecattr_t *vsap, int flag, cred_t *cr)
715 {
716     aclent_t      *aclentp;

718     ASSERT(RW_LOCK_HELD(&ip->i_contents));

720     /* XXX Range check, sanity check, shadow check */
721     /* If an ACL is present, get the data from the shadow inode info */
722     if (ip->i_ufs_acl)
723         return (aclentry2vsecattr(ip->i_ufs_acl, vsap));

725     /*
726      * If no ACLs are present, fabricate one from the mode bits.
727      * This code is almost identical to fs_fab_acl(), but we
728      * already have the mode bits handy, so we'll avoid going
729      * through VOP_GETATTR() again.
730      */

732     vsap->vsa_aclcnt    = 0;
733     vsap->vsa_aclentp   = NULL;
734     vsap->vsa_dfaclcnt  = 0;      /* Default ACLs are not fabricated */
735     vsap->vsa_dfaclentp = NULL;

737     if (vsap->vsa_mask & (VSA_ACLCNT | VSA_ACL))
738         vsap->vsa_aclcnt = 4; /* USER, GROUP, OTHER, and CLASS */

740     if (vsap->vsa_mask & VSA_ACL) {
741         vsap->vsa_aclentp = kmem_zalloc(4 * sizeof (aclent_t),
742             KM_SLEEP);

743         if (vsap->vsa_aclentp == NULL)
744             return (ENOMEM);
744         aclentp = vsap->vsa_aclentp;

746         /* Owner */
747         aclentp->a_type = USER_OBJ;
748         aclentp->a_perm = ((ushort_t)(ip->i_mode & 0700)) >> 6;
749         aclentp->a_id = ip->i_uid; /* Really undefined */
750         aclentp++;

752         /* Group */
753         aclentp->a_type = GROUP_OBJ;
754         aclentp->a_perm = ((ushort_t)(ip->i_mode & 0070)) >> 3;
755         aclentp->a_id = ip->i_gid; /* Really undefined */
756         aclentp++;

758         /* Other */
759         aclentp->a_type = OTHER_OBJ;
760         aclentp->a_perm = ip->i_mode & 0007;
761         aclentp->a_id = 0; /* Really undefined */
762         aclentp++;

764         /* Class */
765         aclentp->a_type = CLASS_OBJ;
766         aclentp->a_perm = ((ushort_t)(ip->i_mode & 0070)) >> 3;
767         aclentp->a_id = 0; /* Really undefined */
```

new/usr/src/uts/common/fs/ufs/ufs_acl.c

2

```
768         ksort((caddr_t)vsap->vsa_aclentp, vsap->vsa_aclcnt,
769             sizeof (aclent_t), cmp2acls);
770     }

772     return (0);
773 }
unchanged_portion_omitted
```


new/usr/src/uts/common/io/aac/aac.c

1

242489 Thu Oct 23 11:04:53 2014

new/usr/src/uts/common/io/aac/aac.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

_____unchanged_portion_omitted_____

```
6189 static int
6190 aac_create_slots(struct aac_softcstate *softs)
6191 {
6192     int i;

6194     softs->total_slots = softs->aac_max_fibs;
6195     softs->io_slot = kmem_zalloc(sizeof (struct aac_slot) * \
6196         softs->total_slots, KM_SLEEP);
6197     if (softs->io_slot == NULL) {
6198         AACDB_PRINT(softs, CE_WARN, "Cannot allocate slot");
6199         return (AACERR);
6200     }
6197     for (i = 0; i < softs->total_slots; i++)
6198         softs->io_slot[i].index = i;
6199     softs->free_io_slot_head = NULL;
6200     softs->total_fibs = 0;
6201     return (AACOK);
6202 }
```

_____unchanged_portion_omitted_____

```

*****
12211 Thu Oct 23 11:04:53 2014
new/usr/src/uts/common/io/arn/arn_phy.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

```

```

335 boolean_t
336 ath9k_hw_init_rf(struct ath_hal *ah, int *status)
337 {
338     struct ath_hal_5416 *ahp = AH5416(ah);
339
340     if (!AR_SREV_9280_10_OR_LATER(ah)) {
341
342         ahp->ah_analogBank0Data =
343             kmem_zalloc((sizeof (uint32_t) *
344             ahp->ah_iniBank0.ia_rows), KM_SLEEP);
345         ahp->ah_analogBank1Data =
346             kmem_zalloc((sizeof (uint32_t) *
347             ahp->ah_iniBank1.ia_rows), KM_SLEEP);
348         ahp->ah_analogBank2Data =
349             kmem_zalloc((sizeof (uint32_t) *
350             ahp->ah_iniBank2.ia_rows), KM_SLEEP);
351         ahp->ah_analogBank3Data =
352             kmem_zalloc((sizeof (uint32_t) *
353             ahp->ah_iniBank3.ia_rows), KM_SLEEP);
354         ahp->ah_analogBank6Data =
355             kmem_zalloc((sizeof (uint32_t) *
356             ahp->ah_iniBank6.ia_rows), KM_SLEEP);
357         ahp->ah_analogBank6TPCData =
358             kmem_zalloc((sizeof (uint32_t) *
359             ahp->ah_iniBank6TPC.ia_rows), KM_SLEEP);
360         ahp->ah_analogBank7Data =
361             kmem_zalloc((sizeof (uint32_t) *
362             ahp->ah_iniBank7.ia_rows), KM_SLEEP);
363
364         if (ahp->ah_analogBank0Data == NULL ||
365             ahp->ah_analogBank1Data == NULL ||
366             ahp->ah_analogBank2Data == NULL ||
367             ahp->ah_analogBank3Data == NULL ||
368             ahp->ah_analogBank6Data == NULL ||
369             ahp->ah_analogBank6TPCData == NULL ||
370             ahp->ah_analogBank7Data == NULL) {
371             ARN_DBG((ARN_DBG_FATAL, "arn: ath9k_hw_init_rf(): "
372             "cannot allocate RF banks\n"));
373             *status = ENOMEM;
374             return (B_FALSE);
375         }
376
377         ahp->ah_addac5416_21 =
378             kmem_zalloc((sizeof (uint32_t) *
379             ahp->ah_iniAddac.ia_rows *
380             ahp->ah_iniAddac.ia_columns), KM_SLEEP);
381         if (ahp->ah_addac5416_21 == NULL) {
382             ARN_DBG((ARN_DBG_FATAL, "arn: ath9k_hw_init_rf(): "
383             "cannot allocate ah_addac5416_21\n"));
384             *status = ENOMEM;
385             return (B_FALSE);
386         }
387
388         ahp->ah_bank6Temp =
389             kmem_zalloc((sizeof (uint32_t) *
390             ahp->ah_iniBank6.ia_rows), KM_SLEEP);
391         if (ahp->ah_bank6Temp == NULL) {
392             ARN_DBG((ARN_DBG_FATAL, "arn: ath9k_hw_init_rf(): "

```

```

393             "cannot allocate ah_bank6Temp\n"));
394             *status = ENOMEM;
395             return (B_FALSE);
396         }
397     }
398
399     return (B_TRUE);
400 }
_____unchanged_portion_omitted_____

```

```

*****
10755 Thu Oct 23 11:04:54 2014
new/usr/src/uts/common/io/beep.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */
21 /*
22  * Copyright 2007 Sun Microsystems, Inc. All rights reserved.
23  * Use is subject to license terms.
24  */

26 #pragma ident      "%Z%M% %I%      %E% SMI"

26 /*
27  * This is the Beep module for supporting keyboard beep for keyboards
28  * that do not have the beeping feature within themselves
29  */
30 */

32 #include <sys/types.h>
33 #include <sys/conf.h>

35 #include <sys/ddi.h>
36 #include <sys/sunddi.h>
37 #include <sys/modctl.h>
38 #include <sys/ddi_impldefs.h>
39 #include <sys/kmem.h>

41 #include <sys/beep.h>
42 #include <sys/inttypes.h>

44 /*
45  * Debug stuff
46  * BEEP_DEBUG used for errors
47  * BEEP_DEBUG1 prints when beep_debug > 1 and used for normal messages
48  */
49 #ifdef DEBUG
50 int beep_debug = 0;
51 #define BEEP_DEBUG(args)      if (beep_debug) cmn_err args
52 #define BEEP_DEBUG1(args)    if (beep_debug > 1) cmn_err args
53 #else
54 #define BEEP_DEBUG(args)
55 #define BEEP_DEBUG1(args)
56 #endif

58 int beep_queue_size = BEEP_QUEUE_SIZE;

```

```

60 /*
61  * Note that mutex_init is not called on the mutex in beep_state,
62  * But assumes that zeroed memory does not need to call mutex_init,
63  * as documented in mutex.c
64  */

66 beep_state_t beep_state;

68 beep_params_t beep_params[] = {
69     {BEEP_CONSOLE, 900, 200},
70     {BEEP_TYPE4, 2000, 0},
71     {BEEP_DEFAULT, 1000, 200}, /* Must be last */
72 };

75 /*
76  * beep_init:
77  * Allocate the beep_queue structure
78  * Initialize beep_state structure
79  * Called from beep driver attach routine
80  */

82 int
83 beep_init(void *arg,
84           beep_on_func_t beep_on_func,
85           beep_off_func_t beep_off_func,
86           beep_freq_func_t beep_freq_func)
87 {
88     beep_entry_t *queue;

90     BEEP_DEBUG1((CE_CONT,
91                "beep_init(0x%lx, 0x%lx, 0x%lx, 0x%lx) : start.",
92                (unsigned long) arg,
93                (unsigned long) beep_on_func,
94                (unsigned long) beep_off_func,
95                (unsigned long) beep_freq_func));

97     mutex_enter(&beep_state.mutex);

99     if (beep_state.mode != BEEP_UNINIT) {
100         mutex_exit(&beep_state.mutex);
101         BEEP_DEBUG((CE_WARN,
102                   "beep_init : beep_state already initialized.));
103         return (DDI_SUCCESS);
104     }

106     queue = kmem_zalloc(sizeof (beep_entry_t) * beep_queue_size,
107                        KM_SLEEP);

109     BEEP_DEBUG1((CE_CONT,
110                "beep_init : beep_queue kmem_zalloc(%d) = 0x%lx.",
111                (int)sizeof (beep_entry_t) * beep_queue_size,
112                (unsigned long)queue));

116     if (queue == NULL) {
117         BEEP_DEBUG((CE_WARN,
118                   "beep_init : kmem_zalloc of beep_queue failed.));
119         return (DDI_FAILURE);
120     }

114     beep_state.arg = arg;
115     beep_state.mode = BEEP_OFF;
116     beep_state.beep_freq = beep_freq_func;
117     beep_state.beep_on = beep_on_func;
118     beep_state.beep_off = beep_off_func;

```

new/usr/src/uts/common/io/beep.c

3

```
119         beep_state.timeout_id = 0;
121         beep_state.queue_head = 0;
122         beep_state.queue_tail = 0;
123         beep_state.queue_size = beep_queue_size;
124         beep_state.queue = queue;
126         mutex_exit(&beep_state.mutex);
128         BEEP_DEBUG1((CE_CONT, "beep_init : done.));
129         return (DDI_SUCCESS);
130 }
```

unchanged_portion_omitted

39281 Thu Oct 23 11:04:54 2014

new/usr/src/uts/common/io/blkdev/blkdev.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

unchanged_portion_omitted

```
1475 static int
1476 bd_flush_write_cache(bd_t *bd, struct dk_callback *dkc)
1477 {
1478     buf_t      *bp;
1479     struct dk_callback *dc;
1480     bd_xfer_impl_t *xi;
1481     int         rv;
1482
1483     if (bd->d_ops.o_sync_cache == NULL) {
1484         return (ENOTSUP);
1485     }
1486     bp = getrbuf(KM_SLEEP);
1487     if ((bp = getrbuf(KM_SLEEP)) == NULL) {
1488         return (ENOMEM);
1489     }
1490     bp->b_resid = 0;
1491     bp->b_bcount = 0;
1492
1493     xi = bd_xfer_alloc(bd, bp, bd->d_ops.o_sync_cache, KM_SLEEP);
1494     if (xi == NULL) {
1495         rv = geterror(bp);
1496         freerbuf(bp);
1497         return (rv);
1498     }
1499
1500     /* Make an asynchronous flush, but only if there is a callback */
1501     if (dkc != NULL && dkc->dkc_callback != NULL) {
1502         /* Make a private copy of the callback structure */
1503         dc = kmem_alloc(sizeof (*dc), KM_SLEEP);
1504         *dc = *dkc;
1505         bp->b_private = dc;
1506         bp->b_iodone = bd_flush_write_cache_done;
1507
1508         bd_submit(bd, xi);
1509         return (0);
1510     }
1511
1512     /* In case there is no callback, perform a synchronous flush */
1513     bd_submit(bd, xi);
1514     (void) biowait(bp);
1515     rv = geterror(bp);
1516     freerbuf(bp);
1517
1518     return (rv);
1519 }
1520
1521 unchanged_portion_omitted
```

new/usr/src/uts/common/io/comstar/lu/stmf_sbd/sbd.c

1

97023 Thu Oct 23 11:04:54 2014

new/usr/src/uts/common/io/comstar/lu/stmf_sbd/sbd.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

_____unchanged_portion_omitted_____

```
3360 /*
3361  * this function creates a local metadata zvol property
3362  */
3363 sbd_status_t
3364 sbd_create_zfs_meta_object(sbd_lu_t *sl)
3365 {
3366     /*
3367      * -allocate 1/2 the property size, the zfs property
3368      * is 8k in size and stored as ascii hex string, all
3369      * we needed is 4k buffer to store the binary data.
3370      * -initialize reader/write lock
3371      */
3372     sl->sl_zfs_meta = kmem_zalloc(ZAP_MAXVALUELEN / 2, KM_SLEEP);
3373     if ((sl->sl_zfs_meta = kmem_zalloc(ZAP_MAXVALUELEN / 2, KM_SLEEP))
3374         == NULL)
3375         return (SBD_FAILURE);
3376     rw_init(&sl->sl_zfs_meta_lock, NULL, RW_DRIVER, NULL);
3377     return (SBD_SUCCESS);
3378 }
3379 _____unchanged_portion_omitted_____
```

```

*****
35222 Thu Oct 23 11:04:54 2014
new/usr/src/uts/common/io/comstar/port/pppt/pppt.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
unchanged_portion_omitted

342 /* ARGSUSED */
343 static int
344 pppt_drv_ioctl(dev_t drv, int cmd, intptr_t argp, int flag, cred_t *cred,
345               int *retval)
346 {
347     int                rc;
348     void               *buf;
349     size_t             buf_size;
350     pppt_iocdata_t     iocd;
351     door_handle_t      new_handle;

353     if (drv_priv(cred) != 0) {
354         return (EPERM);
355     }

357     rc = ddi_copyin((void *)argp, &iocd, sizeof (iocd), flag);
358     if (rc)
359         return (EFAULT);

361     if (iocd.pppt_version != PPPT_VERSION_1)
362         return (EINVAL);

364     switch (cmd) {
365     case PPPT_MESSAGE:

367         /* XXX limit buf_size ? */
368         buf_size = (size_t)iocd.pppt_buf_size;
369         buf = kmem_alloc(buf_size, KM_SLEEP);
370         if (buf == NULL)
371             return (ENOMEM);

371         rc = ddi_copyin((void *)(&iocd.pppt_buf),
372                       buf, buf_size, flag);
373         if (rc) {
374             kmem_free(buf, buf_size);
375             return (EFAULT);
376         }

378         stmf_ic_rx_msg(buf, buf_size);

380         kmem_free(buf, buf_size);
381         break;
382     case PPPT_INSTALL_DOOR:

384         new_handle = door_ki_lookup((int)iocd.pppt_door_fd);
385         if (new_handle == NULL)
386             return (EINVAL);

388         mutex_enter(&pppt_global.global_door_lock);
389         ASSERT(pppt_global.global_svc_state == PSS_ENABLED);
390         if (pppt_global.global_door != NULL) {
391             /*
392              * There can only be one door installed
393              */
394             mutex_exit(&pppt_global.global_door_lock);
395             door_ki_rele(new_handle);
396             return (EBUSY);
397         }

```

```

398         pppt_global.global_door = new_handle;
399         mutex_exit(&pppt_global.global_door_lock);
400         break;
401     }

403     return (rc);
404 }
unchanged_portion_omitted

```

new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_init.c

1

```
*****
116791 Thu Oct 23 11:04:54 2014
new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_init.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

3225 /*
3226 * ql_configure_device_d_id
3227 * Updates device loop ID.
3228 * Also adds to device queue any new devices found on private loop.
3229 *
3230 * Input:
3231 * ha = adapter state pointer.
3232 *
3233 * Returns:
3234 * ql local function return status code.
3235 *
3236 * Context:
3237 * Kernel context.
3238 */
3239 static int
3240 ql_configure_device_d_id(ql_adapter_state_t *ha)
3241 {
3242     port_id_t      d_id;
3243     ql_link_t      *link;
3244     int             rval;
3245     int             loop;
3246     ql_tgt_t       *tq;
3247     ql_dev_id_list_t *list;
3248     uint32_t        list_size;
3249     uint16_t        index, loop_id;
3250     ql_mbx_data_t  mr;
3251     uint8_t         retries = MAX_DEVICE_LOST_RETRY;

3253     QL_PRINT_3(CE_CONT, "%d): started\n", ha->instance);

3255     list_size = sizeof(ql_dev_id_list_t) * DEVICE_LIST_ENTRIES;
3256     list = kmem_zalloc(list_size, KM_SLEEP);
3257     if (list == NULL) {
3258         rval = QL_MEMORY_ALLOC_FAILED;
3259         EL(ha, "failed, rval = %x\n", rval);
3260         return (rval);
3261     }

3258     do {
3259         /*
3260          * Get data from RISC code d_id list to init each device queue.
3261          */
3262         rval = ql_get_id_list(ha, (caddr_t)list, list_size, &mr);
3263         if (rval != QL_SUCCESS) {
3264             kmem_free(list, list_size);
3265             EL(ha, "failed, rval = %x\n", rval);
3266             return (rval);
3267         }

3269         /* Acquire adapter state lock. */
3270         ADAPTER_STATE_LOCK(ha);

3272         /* Mark all queues as unusable. */
3273         for (index = 0; index < DEVICE_HEAD_LIST_SIZE; index++) {
3274             for (link = ha->dev[index].first; link != NULL;
3275                  link = link->nnext) {
3276                 tq = link->base_address;
3277                 DEVICE_QUEUE_LOCK(tq);
```

new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_init.c

2

```
3278         if (!(tq->flags & TQF_PLOGI_PROGRS) &&
3279             !(ha->topology & QL_N_PORT)) {
3280             tq->loop_id = (uint16_t)
3281                 (tq->loop_id | PORT_LOST_ID);
3282         }
3283         DEVICE_QUEUE_UNLOCK(tq);
3284     }
3285 }

3287 /* If device not in queues add new queue. */
3288 for (index = 0; index < mr.mb[1]; index++) {
3289     ql_dev_list(ha, list, index, &d_id, &loop_id);

3291     if (VALID_DEVICE_ID(ha, loop_id)) {
3292         tq = ql_dev_init(ha, d_id, loop_id);
3293         if (tq != NULL) {
3294             tq->loop_id = loop_id;

3296             /* Test for fabric device. */
3297             if (d_id.b.domain !=
3298                 ha->d_id.b.domain ||
3299                 d_id.b.area != ha->d_id.b.area) {
3300                 tq->flags |= TQF_FABRIC_DEVICE;
3301             }

3303             ADAPTER_STATE_UNLOCK(ha);
3304             if (ql_get_port_database(ha, tq,
3305                                     PDF_NONE) == QL_SUCCESS) {
3306                 ADAPTER_STATE_LOCK(ha);
3307                 tq->loop_id = (uint16_t)
3308                     (tq->loop_id &
3309                      ~PORT_LOST_ID);
3310             } else {
3311                 ADAPTER_STATE_LOCK(ha);
3312             }
3313         }
3314     }
3315 }

3317 /* 24xx does not report switch devices in ID list. */
3318 if ((CFG_IST(ha, CFG_CTRL_24258081) &&
3319     ha->topology & (QL_F_PORT | QL_FL_PORT)) {
3320     d_id.b24 = 0xfffffe;
3321     tq = ql_dev_init(ha, d_id, FL_PORT_24XX_HDL);
3322     if (tq != NULL) {
3323         tq->flags |= TQF_FABRIC_DEVICE;
3324         ADAPTER_STATE_UNLOCK(ha);
3325         (void) ql_get_port_database(ha, tq, PDF_NONE);
3326         ADAPTER_STATE_LOCK(ha);
3327     }
3328     d_id.b24 = 0xfffffc;
3329     tq = ql_dev_init(ha, d_id, SNS_24XX_HDL);
3330     if (tq != NULL) {
3331         tq->flags |= TQF_FABRIC_DEVICE;
3332         ADAPTER_STATE_UNLOCK(ha);
3333         if (ha->vp_index != 0) {
3334             (void) ql_login_fport(ha, tq,
3335                                   SNS_24XX_HDL, LFF_NONE, NULL);
3336         }
3337         (void) ql_get_port_database(ha, tq, PDF_NONE);
3338         ADAPTER_STATE_LOCK(ha);
3339     }
3340 }

3342 /* If F_port exists, allocate queue for FL_Port. */
3343 index = ql_alpa_to_index[0xfe];
```



```

3344     d_id.b24 = 0;
3345     if (ha->dev[index].first != NULL) {
3346         tq = ql_dev_init(ha, d_id, (uint16_t)
3347             (CFG_IST(ha, CFG_CTRL_24258081) ?
3348             FL_PORT_24XX_HDL : FL_PORT_LOOP_ID));
3349         if (tq != NULL) {
3350             tq->flags |= TQF_FABRIC_DEVICE;
3351             ADAPTER_STATE_UNLOCK(ha);
3352             (void) ql_get_port_database(ha, tq, PDF_NONE);
3353             ADAPTER_STATE_LOCK(ha);
3354         }
3355     }

3357     /* Allocate queue for broadcast. */
3358     d_id.b24 = 0xffffffff;
3359     (void) ql_dev_init(ha, d_id, (uint16_t)
3360         (CFG_IST(ha, CFG_CTRL_24258081) ? BROADCAST_24XX_HDL :
3361         IP_BROADCAST_LOOP_ID));

3363     /* Check for any devices lost. */
3364     loop = FALSE;
3365     for (index = 0; index < DEVICE_HEAD_LIST_SIZE; index++) {
3366         for (link = ha->dev[index].first; link != NULL;
3367             link = link->next) {
3368             tq = link->base_address;

3370                 if ((tq->loop_id & PORT_LOST_ID) &&
3371                     !(tq->flags & (TQF_INITIATOR_DEVICE |
3372                     TQF_FABRIC_DEVICE))) {
3373                     loop = TRUE;
3374                 }
3375             }
3376         }

3378     /* Release adapter state lock. */
3379     ADAPTER_STATE_UNLOCK(ha);

3381     /* Give devices time to recover. */
3382     if (loop == TRUE) {
3383         drv_usecwait(1000000);
3384     }
3385     } while (retries-- && loop == TRUE &&
3386         !(ha->pha->task_daemon_flags & LOOP_RESYNC_NEEDED));

3388     kmem_free(list, list_size);

3390     if (rval != QL_SUCCESS) {
3391         EL(ha, "failed=%xh\n", rval);
3392     } else {
3393         /*EMPTY*/
3394         QL_PRINT_3(CE_CONT, "(%d): done\n", ha->instance);
3395     }

3397     return (rval);
3398 }

```

unchanged portion omitted

```

4063 /*
4064 * ql_vport_control
4065 * Issue Virtual Port Control command.
4066 *
4067 * Input:
4068 * ha = virtual adapter state pointer.
4069 * cmd = control command.
4070 *
4071 * Returns:

```

```

4072 * ql local function return status code.
4073 *
4074 * Context:
4075 * Kernel context.
4076 */
4077 int
4078 ql_vport_control(ql_adapter_state_t *ha, uint8_t cmd)
4079 {
4080     ql_mbx_ioccb_t *pkt;
4081     uint8_t bit;
4082     int rval;
4083     uint32_t pkt_size;

4085     QL_PRINT_10(CE_CONT, "(%d,%d): started\n", ha->instance, ha->vp_index);

4087     if (ha->vp_index != 0) {
4088         pkt_size = sizeof (ql_mbx_ioccb_t);
4089         pkt = kmem_zalloc(pkt_size, KM_SLEEP);
4090         if (pkt == NULL) {
4091             EL(ha, "failed, kmem_zalloc\n");
4092             return (QL_MEMORY_ALLOC_FAILED);
4093         }

4091         pkt->vpc.entry_type = VP_CONTROL_TYPE;
4092         pkt->vpc.entry_count = 1;
4093         pkt->vpc.command = cmd;
4094         pkt->vpc.vp_count = 1;
4095         bit = (uint8_t)(ha->vp_index - 1);
4096         pkt->vpc.vp_index[bit / 8] = (uint8_t)
4097             (pkt->vpc.vp_index[bit / 8] | BIT_0 << bit % 8);

4099         rval = ql_issue_mbx_ioccb(ha, (caddr_t)pkt, pkt_size);
4100         if (rval == QL_SUCCESS && pkt->vpc.status != 0) {
4101             rval = QL_COMMAND_ERROR;
4102         }

4104         kmem_free(pkt, pkt_size);
4105     } else {
4106         rval = QL_SUCCESS;
4107     }

4109     if (rval != QL_SUCCESS) {
4110         EL(ha, "failed, rval = %xh\n", rval);
4111     } else {
4112         /*EMPTY*/
4113         QL_PRINT_10(CE_CONT, "(%d,%d): done\n", ha->instance,
4114             ha->vp_index);
4115     }
4116     return (rval);
4117 }

4119 /*
4120 * ql_vport_modify
4121 * Issue of Modify Virtual Port command.
4122 *
4123 * Input:
4124 * ha = virtual adapter state pointer.
4125 * cmd = command.
4126 * opt = option.
4127 *
4128 * Context:
4129 * Interrupt or Kernel context, no mailbox commands allowed.
4130 */
4131 int
4132 ql_vport_modify(ql_adapter_state_t *ha, uint8_t cmd, uint8_t opt)
4133 {

```

```
4134     ql_mbx_iocb_t   *pkt;
4135     int              rval;
4136     uint32_t         pkt_size;

4138     QL_PRINT_10(CE_CONT, "(%d,%d): started\n", ha->instance, ha->vp_index);

4140     pkt_size = sizeof (ql_mbx_iocb_t);
4141     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
4142     if (pkt == NULL) {
4143         EL(ha, "failed, kmem_zalloc\n");
4144         return (QL_MEMORY_ALLOC_FAILED);
4145     }

4143     pkt->vpm.entry_type = VP_MODIFY_TYPE;
4144     pkt->vpm.entry_count = 1;
4145     pkt->vpm.command = cmd;
4146     pkt->vpm.vp_count = 1;
4147     pkt->vpm.first_vp_index = ha->vp_index;
4148     pkt->vpm.first_options = opt;
4149     bcopy(ha->loginparams.nport_ww_name.raw_wwn, pkt->vpm.first_port_name,
4150           8);
4151     bcopy(ha->loginparams.node_ww_name.raw_wwn, pkt->vpm.first_node_name,
4152           8);

4154     rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt, pkt_size);
4155     if (rval == QL_SUCCESS && pkt->vpm.status != 0) {
4156         EL(ha, "failed, ql_issue_mbx_iocb=%xh, status=%xh\n", rval,
4157            pkt->vpm.status);
4158         rval = QL_COMMAND_ERROR;
4159     }

4161     kmem_free(pkt, pkt_size);

4163     if (rval != QL_SUCCESS) {
4164         EL(ha, "failed, rval = %xh\n", rval);
4165     } else {
4166         /*EMPTY*/
4167         QL_PRINT_10(CE_CONT, "(%d,%d): done\n", ha->instance,
4168                   ha->vp_index);
4169     }
4170     return (rval);
4171 }
```

unchanged portion omitted

```

*****
52613 Thu Oct 23 11:04:55 2014
new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_ioctl.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

485 /*
486  * Set adapter feature bits in NVRAM
487  */
488 static int
489 ql_set_feature_bits(ql_adapter_state_t *ha, uint16_t features)
490 {
491     int            rval;
492     uint32_t       count;
493     nvram_t        *nv;
494     uint16_t       *wptr;
495     uint8_t        *bptr;
496     uint8_t        csum;
497     uint32_t       start_addr;

499     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

501     if (CFG_IST(ha, CFG_CTRL_24258081)) {
502         EL(ha, "Not supported for 24xx\n");
503         return (EINVAL);
504     }

506     nv = kmem_zalloc(sizeof (*nv), KM_SLEEP);
507     if (nv == NULL) {
508         EL(ha, "failed, kmem_zalloc\n");
509         return (ENOMEM);
510     }

508     rval = ql_lock_nvram(ha, &start_addr, LNF_NVRAM_DATA);
509     if (rval != QL_SUCCESS) {
510         EL(ha, "failed, ql_lock_nvram=%xh\n", rval);
511         kmem_free(nv, sizeof (*nv));
512         return (EIO);
513     }
514     rval = 0;

516     /*
517     * Read off the whole NVRAM
518     */
519     wptr = (uint16_t *)nv;
520     csum = 0;
521     for (count = 0; count < sizeof (nvram_t) / 2; count++) {
522         *wptr = (uint16_t)ql_get_nvram_word(ha, count + start_addr);
523         csum = (uint8_t)(csum + (uint8_t)*wptr);
524         csum = (uint8_t)(csum + (uint8_t)(*wptr >> 8));
525         wptr++;
526     }

528     /*
529     * If the checksum is BAD then fail it right here.
530     */
531     if (csum) {
532         kmem_free(nv, sizeof (*nv));
533         ql_release_nvram(ha);
534         return (EBADF);
535     }

537     nv->adapter_features[0] = (uint8_t)((features & 0xFF00) >> 8);
538     nv->adapter_features[1] = (uint8_t)(features & 0xFF);

```

```

540     /*
541     * Recompute the checksum now
542     */
543     bptr = (uint8_t *)nv;
544     for (count = 0; count < sizeof (nvram_t) - 1; count++) {
545         csum = (uint8_t)(csum + *bptr++);
546     }
547     csum = (uint8_t)(~csum + 1);
548     nv->checksum = csum;

550     /*
551     * Now load the NVRAM
552     */
553     wptr = (uint16_t *)nv;
554     for (count = 0; count < sizeof (nvram_t) / 2; count++) {
555         ql_load_nvram(ha, (uint8_t)(count + start_addr), *wptr++);
556     }

558     /*
559     * Read NVRAM and verify the contents
560     */
561     wptr = (uint16_t *)nv;
562     csum = 0;
563     for (count = 0; count < sizeof (nvram_t) / 2; count++) {
564         if (ql_get_nvram_word(ha, count + start_addr) != *wptr) {
565             rval = EIO;
566             break;
567         }
568         csum = (uint8_t)(csum + (uint8_t)*wptr);
569         csum = (uint8_t)(csum + (uint8_t)(*wptr >> 8));
570         wptr++;
571     }

573     if (csum) {
574         rval = EINVAL;
575     }

577     kmem_free(nv, sizeof (*nv));
578     ql_release_nvram(ha);

580     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

582     return (rval);
583 }

585 /*
586  * Fix this function to update just feature bits and checksum in NVRAM
587  */
588 static int
589 ql_set_nvram_adapter_defaults(ql_adapter_state_t *ha)
590 {
591     int            rval;
592     uint32_t       count;
593     uint32_t       start_addr;

595     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

597     rval = ql_lock_nvram(ha, &start_addr, LNF_NVRAM_DATA);
598     if (rval != QL_SUCCESS) {
599         EL(ha, "failed, ql_lock_nvram=%xh\n", rval);
600         return (EIO);
601     }
602     rval = 0;

604     if (CFG_IST(ha, CFG_CTRL_24258081)) {

```

```

605     nvram_24xx_t    *nv;
606     uint32_t        *longptr;
607     uint32_t        csum = 0;

609     nv = kmem_zalloc(sizeof (*nv), KM_SLEEP);
614     if (nv == NULL) {
615         EL(ha, "failed, kmem_zalloc\n");
616         return (ENOMEM);
617     }

611     nv->nvrाम_version[0] = LSB(ICB_24XX_VERSION);
612     nv->nvrाम_version[1] = MSB(ICB_24XX_VERSION);

614     nv->version[0] = 1;
615     nv->max_frame_length[1] = 8;
616     nv->execution_throttle[0] = 16;
617     nv->login_retry_count[0] = 8;

619     nv->firmware_options_1[0] = BIT_2 | BIT_1;
620     nv->firmware_options_1[1] = BIT_5;
621     nv->firmware_options_2[0] = BIT_5;
622     nv->firmware_options_2[1] = BIT_4;
623     nv->firmware_options_3[1] = BIT_6;

625     /*
626     * Set default host adapter parameters
627     */
628     nv->host_p[0] = BIT_4 | BIT_1;
629     nv->host_p[1] = BIT_3 | BIT_2;
630     nv->reset_delay = 5;
631     nv->max_luns_per_target[0] = 128;
632     nv->port_down_retry_count[0] = 30;
633     nv->link_down_timeout[0] = 30;

635     /*
636     * compute the checksum now
637     */
638     longptr = (uint32_t *)nv;
639     csum = 0;
640     for (count = 0; count < (sizeof (nvram_24xx_t)/4)-1; count++) {
641         csum += *longptr;
642         longptr++;
643     }
644     csum = (uint32_t)(~csum + 1);
645     LITTLE_ENDIAN_32((long)csum);
646     *longptr = csum;

648     /*
649     * Now load the NVRAM
650     */
651     longptr = (uint32_t *)nv;
652     for (count = 0; count < sizeof (nvram_24xx_t) / 4; count++) {
653         (void) ql_24xx_load_nvram(ha,
654             (uint32_t)(count + start_addr), *longptr++);
655     }

657     /*
658     * Read NVRAM and verify the contents
659     */
660     csum = 0;
661     longptr = (uint32_t *)nv;
662     for (count = 0; count < sizeof (nvram_24xx_t) / 4; count++) {
663         rval = ql_24xx_read_flash(ha, count + start_addr,
664             longptr);
665         if (rval != QL_SUCCESS) {
666             EL(ha, "24xx_read_flash failed=%xh\n", rval);

```

```

667         break;
668     }
669     csum += *longptr;
670 }

672     if (csum) {
673         rval = EINVAL;
674     }
675     kmem_free(nv, sizeof (nvram_24xx_t));
676 } else {
677     nvram_t        *nv;
678     uint16_t        *wptr;
679     uint8_t         *bptr;
680     uint8_t         csum;

682     nv = kmem_zalloc(sizeof (*nv), KM_SLEEP);
691     if (nv == NULL) {
692         EL(ha, "failed, kmem_zalloc\n");
693         return (ENOMEM);
694     }
695     /*
696     * Set default initialization control block.
697     */
698     nv->parameter_block_version = ICB_VERSION;
699     nv->firmware_options[0] = BIT_4 | BIT_3 | BIT_2 | BIT_1;
700     nv->firmware_options[1] = BIT_7 | BIT_5 | BIT_2;

702     nv->max_frame_length[1] = 4;
703     nv->max_iocb_allocation[1] = 1;
704     nv->execution_throttle[0] = 16;
705     nv->login_retry_count = 8;
706     nv->port_name[0] = 33;
707     nv->port_name[3] = 224;
708     nv->port_name[4] = 139;
709     nv->login_timeout = 4;

711     /*
712     * Set default host adapter parameters
713     */
714     nv->host_p[0] = BIT_1;
715     nv->host_p[1] = BIT_2;
716     nv->reset_delay = 5;
717     nv->port_down_retry_count = 8;
718     nv->maximum_luns_per_target[0] = 8;

720     /*
721     * compute the checksum now
722     */
723     bptr = (uint8_t *)nv;
724     csum = 0;
725     for (count = 0; count < sizeof (nvram_t) - 1; count++) {
726         csum = (uint8_t)(csum + *bptr++);
727     }
728     csum = (uint8_t)(~csum + 1);
729     nv->checksum = csum;

731     /*
732     * Now load the NVRAM
733     */
734     wptr = (uint16_t *)nv;
735     for (count = 0; count < sizeof (nvram_t) / 2; count++) {
736         ql_load_nvram(ha, (uint8_t)(count + start_addr),
737             *wptr++);
738     }

740     /*

```

```

729     * Read NVRAM and verify the contents
730     */
731     wptr = (uint16_t *)nv;
732     csum = 0;
733     for (count = 0; count < sizeof (nvram_t) / 2; count++) {
734         if (ql_get_nvram_word(ha, count + start_addr) !=
735             *wptr) {
736             rval = EIO;
737             break;
738         }
739         csum = (uint8_t)(csum + (uint8_t)*wptr);
740         csum = (uint8_t)(csum + (uint8_t)(*wptr >> 8));
741         wptr++;
742     }
743     if (csum) {
744         rval = EINVAL;
745     }
746     kmem_free(nv, sizeof (*nv));
747 }
748 ql_release_nvram(ha);

750     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

752     return (rval);
753 }

```

unchanged portion omitted

```

905 /*
906 * ql_nv_util_load
907 * Loads NVRAM from application.
908 *
909 * Input:
910 *     ha = adapter state pointer.
911 *     bp = user buffer address.
912 *
913 * Returns:
914 *
915 * Context:
916 *     Kernel context.
917 */
918 int
919 ql_nv_util_load(ql_adapter_state_t *ha, void *bp, int mode)
920 {
921     uint8_t     cnt;
922     void        *nv;
923     uint16_t    *wptr;
924     uint16_t    data;
925     uint32_t    start_addr, *lptr, data32;
926     nvram_t    *nptr;
927     int         rval;

929     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

931     nv = kmem_zalloc(ha->nvram_cache->size, KM_SLEEP);
932     if ((nv = kmem_zalloc(ha->nvram_cache->size, KM_SLEEP)) == NULL) {
933         EL(ha, "failed, kmem_zalloc\n");
934         return (ENOMEM);
935     }

936     if (ddi_copyin(bp, nv, ha->nvram_cache->size, mode) != 0) {
937         EL(ha, "Buffer copy failed\n");
938         kmem_free(nv, ha->nvram_cache->size);
939         return (EFAULT);
940     }

941     /* See if the buffer passed to us looks sane */

```

```

940     nptr = (nvram_t *)nv;
941     if (nptr->id[0] != 'I' || nptr->id[1] != 'S' || nptr->id[2] != 'P' ||
942         nptr->id[3] != ' ') {
943         EL(ha, "failed, buffer sanity check\n");
944         kmem_free(nv, ha->nvram_cache->size);
945         return (EINVAL);
946     }

948     /* Quiesce I/O */
949     if (ql_stall_driver(ha, 0) != QL_SUCCESS) {
950         EL(ha, "ql_stall_driver failed\n");
951         kmem_free(nv, ha->nvram_cache->size);
952         return (EBUSY);
953     }

955     rval = ql_lock_nvram(ha, &start_addr, LNF_NVRAM_DATA);
956     if (rval != QL_SUCCESS) {
957         EL(ha, "failed, ql_lock_nvram=%xh\n", rval);
958         kmem_free(nv, ha->nvram_cache->size);
959         ql_restart_driver(ha);
960         return (EIO);
961     }

963     /* Load NVRAM. */
964     if (CFG_IST(ha, CFG_CTRL_258081)) {
965         GLOBAL_HW_UNLOCK();
966         start_addr &= ~ha->flash_data_addr;
967         start_addr <= 2;
968         if ((rval = ql_r_m_w_flash(ha, bp, ha->nvram_cache->size,
969             start_addr, mode)) != QL_SUCCESS) {
970             EL(ha, "nvram load failed, rval = %0xh\n", rval);
971         }
972         GLOBAL_HW_LOCK();
973     } else if (CFG_IST(ha, CFG_CTRL_2422)) {
974         lptr = (uint32_t *)nv;
975         for (cnt = 0; cnt < ha->nvram_cache->size / 4; cnt++) {
976             data32 = *lptr++;
977             LITTLE_ENDIAN_32(&data32);
978             rval = ql_24xx_load_nvram(ha, cnt + start_addr,
979                 data32);
980             if (rval != QL_SUCCESS) {
981                 EL(ha, "failed, 24xx_load_nvram=%xh\n", rval);
982                 break;
983             }
984         }
985     } else {
986         wptr = (uint16_t *)nv;
987         for (cnt = 0; cnt < ha->nvram_cache->size / 2; cnt++) {
988             data = *wptr++;
989             LITTLE_ENDIAN_16(&data);
990             ql_load_nvram(ha, (uint8_t)(cnt + start_addr), data);
991         }
992     }
993     /* switch to the new one */
994     NVRAM_CACHE_LOCK(ha);

996     kmem_free(ha->nvram_cache->cache, ha->nvram_cache->size);
997     ha->nvram_cache->cache = (void *)nptr;

999     NVRAM_CACHE_UNLOCK(ha);

1001     ql_release_nvram(ha);
1002     ql_restart_driver(ha);

1004     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

```

```

1006     if (rval == QL_SUCCESS) {
1007         return (0);
1008     }

1010     return (EFAULT);
1011 }
unchanged_portion_omitted

1120 /*
1121  * ql_vpd_load
1122  * Loads VPD from application.
1123  *
1124  * Input:
1125  *   ha = adapter state pointer.
1126  *   bp = user buffer address.
1127  *
1128  * Returns:
1129  *
1130  * Context:
1131  *   Kernel context.
1132  */
1133 int
1134 ql_vpd_load(ql_adapter_state_t *ha, void *bp, int mode)
1135 {
1136     uint8_t      cnt;
1137     uint8_t      *vpd, *vpdptr, *vbuf;
1138     uint32_t     start_addr, vpd_size, *lptr, data32;
1139     int          rval;

1141     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

1143     if ((CFG_IST(ha, CFG_CTRL_24258081)) == 0) {
1144         EL(ha, "unsupported adapter feature\n");
1145         return (ENOTSUP);
1146     }

1148     vpd_size = QL_24XX_VPD_SIZE;

1150     vpd = kmem_zalloc(vpd_size, KM_SLEEP);
1151     if ((vpd = kmem_zalloc(vpd_size, KM_SLEEP)) == NULL) {
1152         EL(ha, "failed, kmem_zalloc\n");
1153         return (ENOMEM);
1154     }

1152     if (ddi_copyin(bp, vpd, vpd_size, mode) != 0) {
1153         EL(ha, "Buffer copy failed\n");
1154         kmem_free(vpd, vpd_size);
1155         return (EFAULT);
1156     }

1158     /* Sanity check the user supplied data via checksum */
1159     if ((vpdptr = ql_vpd_findtag(ha, vpd, "RV")) == NULL) {
1160         EL(ha, "vpd RV tag missing\n");
1161         kmem_free(vpd, vpd_size);
1162         return (EINVAL);
1163     }

1165     vpdptr += 3;
1166     cnt = 0;
1167     vbuf = vpd;
1168     while (vbuf <= vpdptr) {
1169         cnt += *vbuf++;
1170     }
1171     if (cnt != 0) {
1172         EL(ha, "mismatched checksum, cal=%xh, passed=%xh\n",
1173            (uint8_t)cnt, (uintptr_t)vpdptr);

```

```

1174         kmem_free(vpd, vpd_size);
1175         return (EINVAL);
1176     }

1178     /* Quiesce I/O */
1179     if (ql_stall_driver(ha, 0) != QL_SUCCESS) {
1180         EL(ha, "ql_stall_driver failed\n");
1181         kmem_free(vpd, vpd_size);
1182         return (EBUSY);
1183     }

1185     rval = ql_lock_nvram(ha, &start_addr, LNF_VPD_DATA);
1186     if (rval != QL_SUCCESS) {
1187         EL(ha, "failed, ql_lock_nvram=%xh\n", rval);
1188         kmem_free(vpd, vpd_size);
1189         ql_restart_driver(ha);
1190         return (EIO);
1191     }

1193     /* Load VPD. */
1194     if (CFG_IST(ha, CFG_CTRL_258081)) {
1195         GLOBAL_HW_UNLOCK();
1196         start_addr &= ~ha->flash_data_addr;
1197         start_addr <<= 2;
1198         if ((rval = ql_r_m_w_flash(ha, bp, vpd_size, start_addr,
1199             mode)) != QL_SUCCESS) {
1200             EL(ha, "vpd load error: %xh\n", rval);
1201         }
1202         GLOBAL_HW_LOCK();
1203     } else {
1204         lptr = (uint32_t *)vpd;
1205         for (cnt = 0; cnt < vpd_size / 4; cnt++) {
1206             data32 = *lptr++;
1207             LITTLE_ENDIAN_32(&data32);
1208             rval = ql_24xx_load_nvram(ha, cnt + start_addr,
1209                 data32);
1210             if (rval != QL_SUCCESS) {
1211                 EL(ha, "failed, 24xx_load_nvram=%xh\n", rval);
1212                 break;
1213             }
1214         }
1215     }

1217     kmem_free(vpd, vpd_size);

1219     /* Update the vcache */
1220     CACHE_LOCK(ha);

1222     if (rval != QL_SUCCESS) {
1223         EL(ha, "failed, load\n");
1224     } else if ((ha->vcache == NULL) && ((ha->vcache =
1225         kmem_zalloc(vpd_size, KM_SLEEP)) == NULL)) {
1226         EL(ha, "failed, kmem_zalloc2\n");
1227     } else if (ddi_copyin(bp, ha->vcache, vpd_size, mode) != 0) {
1228         EL(ha, "Buffer copy2 failed\n");
1229         kmem_free(ha->vcache, vpd_size);
1230         ha->vcache = NULL;
1231     }

1233     CACHE_UNLOCK(ha);

1235     ql_release_nvram(ha);
1236     ql_restart_driver(ha);

1238     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

```

```

1240     if (rval == QL_SUCCESS) {
1241         return (0);
1242     }
1244     return (EFAULT);
1245 }

1247 /*
1248  * ql_vpd_dump
1249  * Dumps VPD to application buffer.
1250  *
1251  * Input:
1252  *   ha = adapter state pointer.
1253  *   bp = user buffer address.
1254  *
1255  * Returns:
1256  *
1257  * Context:
1258  *   Kernel context.
1259  */
1260 int
1261 ql_vpd_dump(ql_adapter_state_t *ha, void *bp, int mode)
1262 {
1263     uint8_t      cnt;
1264     void         *vpd;
1265     uint32_t     start_addr, vpd_size, *lptr;
1266     int          rval = 0;

1268     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

1270     if ((CFG_IST(ha, CFG_CTRL_24258081)) == 0) {
1271         EL(ha, "unsupported adapter feature\n");
1272         return (EACCES);
1273     }

1275     vpd_size = QL_24XX_VPD_SIZE;

1277     CACHE_LOCK(ha);

1279     if (ha->vcache != NULL) {
1280         /* copy back the vpd cache data */
1281         if (ddi_copyout(ha->vcache, bp, vpd_size, mode) != 0) {
1282             EL(ha, "Buffer copy failed\n");
1283             rval = EFAULT;
1284         }
1285         CACHE_UNLOCK(ha);
1286         return (rval);
1287     }

1289     vpd = kmem_zalloc(vpd_size, KM_SLEEP);
1307     if ((vpd = kmem_zalloc(vpd_size, KM_SLEEP)) == NULL) {
1308         CACHE_UNLOCK(ha);
1309         EL(ha, "failed, kmem_zalloc\n");
1310         return (ENOMEM);
1311     }

1291     /* Quiesce I/O */
1292     if (ql_stall_driver(ha, 0) != QL_SUCCESS) {
1293         CACHE_UNLOCK(ha);
1294         EL(ha, "ql_stall_driver failed\n");
1295         kmem_free(vpd, vpd_size);
1296         return (EBUSY);
1297     }

1299     rval = ql_lock_nvram(ha, &start_addr, LNF_VPD_DATA);
1300     if (rval != QL_SUCCESS) {

```

```

1301         CACHE_UNLOCK(ha);
1302         EL(ha, "failed, ql_lock_nvram=%xh\n", rval);
1303         kmem_free(vpd, vpd_size);
1304         ql_restart_driver(ha);
1305         return (EIO);
1306     }
1308     /* Dump VPD. */
1309     lptr = (uint32_t *)vpd;

1311     for (cnt = 0; cnt < vpd_size / 4; cnt++) {
1312         rval = ql_24xx_read_flash(ha, start_addr++, lptr);
1313         if (rval != QL_SUCCESS) {
1314             EL(ha, "read_flash failed=%xh\n", rval);
1315             rval = EAGAIN;
1316             break;
1317         }
1318         LITTLE_ENDIAN_32(lptr);
1319         lptr++;
1320     }

1322     ql_release_nvram(ha);
1323     ql_restart_driver(ha);

1325     if (ddi_copyout(vpd, bp, vpd_size, mode) != 0) {
1326         CACHE_UNLOCK(ha);
1327         EL(ha, "Buffer copy failed\n");
1328         kmem_free(vpd, vpd_size);
1329         return (EFAULT);
1330     }

1332     ha->vcache = vpd;

1334     CACHE_UNLOCK(ha);

1336     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

1338     if (rval != QL_SUCCESS) {
1339         return (EFAULT);
1340     } else {
1341         return (0);
1342     }
1343 }
    unchanged portion omitted

1525 /*
1526  * ql_r_m_w_flash
1527  * Read modify write from user space to flash.
1528  *
1529  * Input:
1530  *   ha: adapter state pointer.
1531  *   dp: source byte pointer.
1532  *   bc: byte count.
1533  *   faddr: flash byte address.
1534  *   mode: flags.
1535  *
1536  * Returns:
1537  *   ql local function return status code.
1538  *
1539  * Context:
1540  *   Kernel context.
1541  */
1542 int
1543 ql_r_m_w_flash(ql_adapter_state_t *ha, caddr_t dp, uint32_t bc, uint32_t faddr,
1544               int mode)
1545 {

```

```

1546     uint8_t      *bp;
1547     uint32_t     xfer, bsize, saddr, ofst;
1548     int          rval = 0;

1550     QL_PRINT_9(CE_CONT, "(%d): started, dp=%ph, faddr=%xh, bc=%xh\n",
1551             ha->instance, (void *)dp, faddr, bc);

1553     bsize = ha->xiocntl->fdesc.block_size;
1554     saddr = faddr & ~(bsize - 1);
1555     ofst = faddr & (bsize - 1);

1557     bp = kmem_zalloc(bsize, KM_SLEEP);
1579     if ((bp = kmem_zalloc(bsize, KM_SLEEP)) == NULL) {
1580         EL(ha, "kmem_zalloc=null\n");
1581         return (QL_MEMORY_ALLOC_FAILED);
1582     }

1559     while (bc) {
1560         xfer = bc > bsize ? bsize : bc;
1561         if (ofst + xfer > bsize) {
1562             xfer = bsize - ofst;
1563         }
1564         QL_PRINT_9(CE_CONT, "(%d): dp=%ph, saddr=%xh, bc=%xh, "
1565                 "ofst=%xh, xfer=%xh\n", ha->instance, (void *)dp, saddr,
1566                 bc, ofst, xfer);

1568         if (ofst || xfer < bsize) {
1569             /* Dump Flash sector. */
1570             if ((rval = ql_dump_fcode(ha, bp, bsize, saddr)) !=
1571                 QL_SUCCESS) {
1572                 EL(ha, "dump_flash status=%x\n", rval);
1573                 break;
1574             }
1575         }

1577         /* Set new data. */
1578         if ((rval = ddi_copyin(dp, (caddr_t)(bp + ofst), xfer,
1579             mode)) != 0) {
1580             EL(ha, "ddi_copyin status=%xh, dp=%ph, ofst=%xh, "
1581                 "xfer=%xh\n", rval, (void *)dp, ofst, xfer);
1582             rval = QL_FUNCTION_FAILED;
1583             break;
1584         }

1586         /* Write to flash. */
1587         if ((rval = ql_load_fcode(ha, bp, bsize, saddr)) !=
1588             QL_SUCCESS) {
1589             EL(ha, "load_flash status=%x\n", rval);
1590             break;
1591         }
1592         bc -= xfer;
1593         dp += xfer;
1594         saddr += bsize;
1595         ofst = 0;
1596     }

1598     kmem_free(bp, bsize);

1600     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

1602     return (rval);
1603 }
_____unchanged_portion_omitted_____

1990 /*
1991  * ql_adm_prop_update_int

```

```

1992  *     Performs qladm QL_PROP_UPDATE_INT command
1993  *
1994  * Input:
1995  *     ha:     adapter state pointer.
1996  *     dop:    ql_adm_op_t structure pointer.
1997  *     mode:   flags.
1998  *
1999  * Returns:
2000  *
2001  * Context:
2002  *     Kernel context.
2003  */
2004 static int
2005 ql_adm_prop_update_int(ql_adapter_state_t *ha, ql_adm_op_t *dop, int mode)
2006 {
2007     char     *prop_name;
2008     int      rval;

2010     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

2012     prop_name = kmem_zalloc(dop->length, KM_SLEEP);
2038     if (prop_name == NULL) {
2039         EL(ha, "failed, kmem_zalloc\n");
2040         return (ENOMEM);
2041     }

2014     if (ddi_copyin((void *) (uintptr_t) dop->buffer, prop_name, dop->length,
2015         mode) != 0) {
2016         EL(ha, "failed, prop_name ddi_copyin\n");
2017         kmem_free(prop_name, dop->length);
2018         return (EFAULT);
2019     }

2021     /*LINTED [Solaris DDI_DEV_T_ANY Lint warning]*/
2022     if ((rval = ddi_prop_update_int(DDI_DEV_T_NONE, ha->dip, prop_name,
2023         (int) dop->option)) != DDI_PROP_SUCCESS) {
2024         EL(ha, "failed, prop_update=%xh\n", rval);
2025         kmem_free(prop_name, dop->length);
2026         return (EINVAL);
2027     }

2029     kmem_free(prop_name, dop->length);

2031     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

2033     return (0);
2034 }

2036 /*
2037  * ql_adm_fw_dump
2038  *     Performs qladm QL_FW_DUMP command
2039  *
2040  * Input:
2041  *     ha:     adapter state pointer.
2042  *     dop:    ql_adm_op_t structure pointer.
2043  *     udop:   user space ql_adm_op_t structure pointer.
2044  *     mode:   flags.
2045  *
2046  * Returns:
2047  *
2048  * Context:
2049  *     Kernel context.
2050  */
2051 static int
2052 ql_adm_fw_dump(ql_adapter_state_t *ha, ql_adm_op_t *dop, void *udop, int mode)
2053 {

```



```

2054     caddr_t dmp;
2056     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);
2058     if (dop->length < ha->risc_dump_size) {
2059         EL(ha, "failed, incorrect length=%xh, size=%xh\n",
2060            dop->length, ha->risc_dump_size);
2061         return (EINVAL);
2062     }
2064     if (ha->ql_dump_state & QL_DUMP_VALID) {
2065         dmp = kmem_zalloc(ha->risc_dump_size, KM_SLEEP);
2066         if (dmp == NULL) {
2067             EL(ha, "failed, kmem_zalloc\n");
2068             return (ENOMEM);
2069         }
2071         dop->length = (uint32_t)ql_ascii_fw_dump(ha, dmp);
2072         if (ddi_copyout((void *)dmp, (void *)dop->buffer,
2073            dop->length, mode) != 0) {
2074             EL(ha, "failed, ddi_copyout\n");
2075             kmem_free(dmp, ha->risc_dump_size);
2076             return (EFAULT);
2077         }
2078         kmem_free(dmp, ha->risc_dump_size);
2079         ha->ql_dump_state |= QL_DUMP_UPLOADED;
2080     } else {
2081         EL(ha, "failed, no dump file\n");
2082         dop->length = 0;
2083     }
2084     if (ddi_copyout(dop, udop, sizeof(ql_adm_op_t), mode) != 0) {
2085         EL(ha, "failed, driver_op_t ddi_copyout\n");
2086         return (EFAULT);
2087     }
2088     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
2089     return (0);
2090 }

```

unchanged portion omitted

```

2169 /*
2170 * ql_adm_flash_load
2171 * Performs qladm QL_FLASH_LOAD command
2172 *
2173 * Input:
2174 *   ha: adapter state pointer.
2175 *   dop: ql_adm_op_t structure pointer.
2176 *   mode: flags.
2177 *
2178 * Returns:
2179 *
2180 * Context:
2181 *   Kernel context.
2182 */
2183 static int
2184 ql_adm_flash_load(ql_adapter_state_t *ha, ql_adm_op_t *dop, int mode)
2185 {
2186     uint8_t *dp;
2187     int rval;
2188 }
2189     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);

```

```

2191     dp = kmem_zalloc(dop->length, KM_SLEEP);
2192     if ((dp = kmem_zalloc(dop->length, KM_SLEEP)) == NULL) {
2193         EL(ha, "failed, kmem_zalloc\n");
2194         return (ENOMEM);
2195     }
2196     if (ddi_copyin((void *)dop->buffer, dp, dop->length,
2197        mode) != 0) {
2198         EL(ha, "ddi_copyin failed\n");
2199         kmem_free(dp, dop->length);
2200         return (EFAULT);
2201     }
2202     if (ql_stall_driver(ha, 0) != QL_SUCCESS) {
2203         EL(ha, "ql_stall_driver failed\n");
2204         kmem_free(dp, dop->length);
2205         return (EBUSY);
2206     }
2207     rval = (CFG_IST(ha, CFG_CTRL_24258081) ?
2208        ql_24xx_load_flash(ha, dp, dop->length, dop->option) :
2209        ql_load_flash(ha, dp, dop->length));
2210     ql_restart_driver(ha);
2211     kmem_free(dp, dop->length);
2212     if (rval != QL_SUCCESS) {
2213         EL(ha, "failed\n");
2214         return (EIO);
2215     }
2216     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
2217     return (0);
2218 }

```

unchanged portion omitted

new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_mbx.c

1

```
*****
112730 Thu Oct 23 11:04:55 2014
new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_mbx.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
unchanged_portion_omitted_

1251 /*
1252 * ql_task_mgmt_iocb
1253 *     Function issues task management IOCB.
1254 *
1255 * Input:
1256 *     ha:     adapter state pointer.
1257 *     tq:     target queue pointer.
1258 *     lun:    LUN.
1259 *     flags:  control flags.
1260 *     delay:  seconds.
1261 *
1262 * Returns:
1263 *     ql local function return status code.
1264 *
1265 * Context:
1266 *     Kernel context
1267 */
1268 static int
1269 ql_task_mgmt_iocb(ql_adapter_state_t *ha, ql_tgt_t *tq, uint16_t lun,
1270                 uint32_t flags, uint16_t delay)
1271 {
1272     ql_mbx_iocb_t *pkt;
1273     int rval;
1274     uint32_t pkt_size;

1276     QL_PRINT_3(CE_CONT, "(%d): started\n", ha->instance);

1278     pkt_size = sizeof(ql_mbx_iocb_t);
1279     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
1280     if (pkt == NULL) {
1281         EL(ha, "failed, kmem_zalloc\n");
1282         return (QL_MEMORY_ALLOC_FAILED);
1283     }

1281     pkt->mgmt.entry_type = TASK_MGMT_TYPE;
1282     pkt->mgmt.entry_count = 1;

1284     pkt->mgmt.n_port_hdl = (uint16_t)LE_16(tq->loop_id);
1285     pkt->mgmt.delay = (uint16_t)LE_16(delay);
1286     pkt->mgmt.timeout = LE_16(MAILBOX_TOV);
1287     pkt->mgmt.fcp_lun[2] = LSB(lun);
1288     pkt->mgmt.fcp_lun[3] = MSB(lun);
1289     pkt->mgmt.control_flags = LE_32(flags);
1290     pkt->mgmt.target_id[0] = tq->d_id.b.al_pa;
1291     pkt->mgmt.target_id[1] = tq->d_id.b.area;
1292     pkt->mgmt.target_id[2] = tq->d_id.b.domain;
1293     pkt->mgmt.vp_index = ha->vp_index;

1295     rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt, pkt_size);
1296     if (rval == QL_SUCCESS && (pkt->sts24.entry_status & 0x3c) != 0) {
1297         EL(ha, "failed, entry_status=%xh, d_id=%xh\n",
1298            pkt->sts24.entry_status, tq->d_id.b24);
1299         rval = QL_FUNCTION_PARAMETER_ERROR;
1300     }

1302     LITTLE_ENDIAN_16(&pkt->sts24.comp_status);

1304     if (rval == QL_SUCCESS && pkt->sts24.comp_status != CS_COMPLETE) {
```

new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_mbx.c

2

```
1305         EL(ha, "failed, comp_status=%xh, d_id=%xh\n",
1306            pkt->sts24.comp_status, tq->d_id.b24);
1307         rval = QL_FUNCTION_FAILED;
1308     }

1310     kmem_free(pkt, pkt_size);

1312     if (rval != QL_SUCCESS) {
1313         EL(ha, "failed, rval = %xh\n", rval);
1314     } else {
1315         /*EMPTY*/
1316         QL_PRINT_3(CE_CONT, "(%d): done\n", ha->instance);
1317     }

1319     return (rval);
1320 }
unchanged_portion_omitted_

1617 /*
1618 * ql_log_iocb
1619 *     Function issues login/logout IOCB.
1620 *
1621 * Input:
1622 *     ha:     adapter state pointer.
1623 *     tq:     target queue pointer.
1624 *     loop_id: FC Loop ID.
1625 *     flags:  control flags.
1626 *     mr:     pointer for mailbox data.
1627 *
1628 * Returns:
1629 *     ql local function return status code.
1630 *
1631 * Context:
1632 *     Kernel context.
1633 */
1634 int
1635 ql_log_iocb(ql_adapter_state_t *ha, ql_tgt_t *tq, uint16_t loop_id,
1636            uint16_t flags, ql_mbx_data_t *mr)
1637 {
1638     ql_mbx_iocb_t *pkt;
1639     int rval;
1640     uint32_t pkt_size;

1642     QL_PRINT_3(CE_CONT, "(%d): started\n", ha->instance);

1644     pkt_size = sizeof(ql_mbx_iocb_t);
1645     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
1646     if (pkt == NULL) {
1647         EL(ha, "failed, kmem_zalloc\n");
1648         return (QL_MEMORY_ALLOC_FAILED);
1649     }

1647     pkt->log.entry_type = LOG_TYPE;
1648     pkt->log.entry_count = 1;
1649     pkt->log.n_port_hdl = (uint16_t)LE_16(loop_id);
1650     pkt->log.control_flags = (uint16_t)LE_16(flags);
1651     pkt->log.port_id[0] = tq->d_id.b.al_pa;
1652     pkt->log.port_id[1] = tq->d_id.b.area;
1653     pkt->log.port_id[2] = tq->d_id.b.domain;
1654     pkt->log.vp_index = ha->vp_index;

1656     rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt, pkt_size);
1657     if (rval == QL_SUCCESS && (pkt->log.entry_status & 0x3c) != 0) {
1658         EL(ha, "failed, entry_status=%xh, d_id=%xh\n",
1659            pkt->log.entry_status, tq->d_id.b24);
1660         rval = QL_FUNCTION_PARAMETER_ERROR;
1661     }
```

```

1661     }
1663     if (rval == QL_SUCCESS) {
1664         if (pkt->log.rsp_size == 0xB) {
1665             LITTLE_ENDIAN_32(&pkt->log.io_param[5]);
1666             tq->cmn_features = MSW(pkt->log.io_param[5]);
1667             LITTLE_ENDIAN_32(&pkt->log.io_param[6]);
1668             tq->conc_sequences = MSW(pkt->log.io_param[6]);
1669             tq->relative_offset = LSW(pkt->log.io_param[6]);
1670             LITTLE_ENDIAN_32(&pkt->log.io_param[9]);
1671             tq->class3_recipient_ctl = MSW(pkt->log.io_param[9]);
1672             tq->class3_conc_sequences = LSW(pkt->log.io_param[9]);
1673             LITTLE_ENDIAN_32(&pkt->log.io_param[10]);
1674             tq->class3_open_sequences_per_exch =
1675                 MSW(pkt->log.io_param[10]);
1676             tq->prli_payload_length = 0x14;
1677         }
1678         if (mr != NULL) {
1679             LITTLE_ENDIAN_16(&pkt->log.status);
1680             LITTLE_ENDIAN_32(&pkt->log.io_param[0]);
1681             LITTLE_ENDIAN_32(&pkt->log.io_param[1]);
1683
1684             if (pkt->log.status != CS_COMPLETE) {
1685                 EL(ha, "failed, status=%xh, iop0=%xh, iop1="
1686                     "%xh\n", pkt->log.status,
1687                     pkt->log.io_param[0],
1688                     pkt->log.io_param[1]);
1689
1690                 switch (pkt->log.io_param[0]) {
1691                     case CS0_NO_LINK:
1692                     case CS0_FIRMWARE_NOT_READY:
1693                         mr->mb[0] = MBS_COMMAND_ERROR;
1694                         mr->mb[1] = 1;
1695                         break;
1696                     case CS0_NO_IOCB:
1697                     case CS0_NO_PCB_ALLOCATED:
1698                         mr->mb[0] = MBS_COMMAND_ERROR;
1699                         mr->mb[1] = 2;
1700                         break;
1701                     case CS0_NO_EXCH_CTRL_BLK:
1702                         mr->mb[0] = MBS_COMMAND_ERROR;
1703                         mr->mb[1] = 3;
1704                         break;
1705                     case CS0_COMMAND_FAILED:
1706                         mr->mb[0] = MBS_COMMAND_ERROR;
1707                         mr->mb[1] = 4;
1708                         switch (LSB(pkt->log.io_param[1])) {
1709                             case CS1_PLOGI_RESPONSE_FAILED:
1710                                 mr->mb[2] = 3;
1711                                 break;
1712                             case CS1_PRLI_FAILED:
1713                                 mr->mb[2] = 4;
1714                                 break;
1715                             case CS1_PRLI_RESPONSE_FAILED:
1716                                 mr->mb[2] = 5;
1717                                 break;
1718                             case CS1_COMMAND_LOGGED_OUT:
1719                                 mr->mb[2] = 7;
1720                                 break;
1721                             case CS1_PLOGI_FAILED:
1722                             default:
1723                                 EL(ha, "log iop1 = %xh\n",
1724                                     LSB(pkt->log.io_param[1]))
1725                                 mr->mb[2] = 2;
1726                                 break;
1727                         }
1728                 }
1729             }

```

```

1727         break;
1728     case CS0_PORT_NOT_LOGGED_IN:
1729         mr->mb[0] = MBS_COMMAND_ERROR;
1730         mr->mb[1] = 4;
1731         mr->mb[2] = 7;
1732         break;
1733     case CS0_NO_FLOGI_ACC:
1734     case CS0_NO_FABRIC_PRESENT:
1735         mr->mb[0] = MBS_COMMAND_ERROR;
1736         mr->mb[1] = 5;
1737         break;
1738     case CS0_ELS_REJECT_RECEIVED:
1739         mr->mb[0] = MBS_COMMAND_ERROR;
1740         mr->mb[1] = 0xd;
1741         break;
1742     case CS0_PORT_ID_USED:
1743         mr->mb[0] = MBS_PORT_ID_USED;
1744         mr->mb[1] = LSW(pkt->log.io_param[1]);
1745         break;
1746     case CS0_N_PORT_HANDLE_USED:
1747         mr->mb[0] = MBS_LOOP_ID_USED;
1748         mr->mb[1] = MSW(pkt->log.io_param[1]);
1749         mr->mb[2] = LSW(pkt->log.io_param[1]);
1750         break;
1751     case CS0_NO_N_PORT_HANDLE_AVAILABLE:
1752         mr->mb[0] = MBS_ALL_IDS_IN_USE;
1753         break;
1754     case CS0_CMD_PARAMETER_ERROR:
1755     default:
1756         EL(ha, "pkt->log iop0=%xh\n",
1757             pkt->log.io_param[0]);
1758         mr->mb[0] =
1759             MBS_COMMAND_PARAMETER_ERROR;
1760         break;
1761     } else {
1762         QL_PRINT_3(CE_CONT, "(%d): status=%xh\n",
1763             ha->instance, pkt->log.status);
1764
1765         mr->mb[0] = MBS_COMMAND_COMPLETE;
1766         mr->mb[1] = (uint16_t)
1767             (pkt->log.io_param[0] & BIT_4 ? 0 : BIT_0);
1768         if (pkt->log.io_param[0] & BIT_8) {
1769             mr->mb[1] = (uint16_t)
1770                 (mr->mb[1] | BIT_1);
1771         }
1772         rval = mr->mb[0];
1773     }
1774 }
1775
1777 }
1779
1779     kmem_free(pkt, pkt_size);
1781
1781     if (rval != QL_SUCCESS) {
1782         EL(ha, "failed=%xh, d_id=%xh\n", rval, tq->d_id.b24);
1783     } else {
1784         /*EMPTY*/
1785         QL_PRINT_3(CE_CONT, "(%d): done\n", ha->instance);
1786     }
1788     return (rval);
1789 }

```

unchanged_portion_omitted

```

2664 /*

```

```

2665 * ql_abort_cmd_iocb
2666 *     Function issues abort command IOCB.
2667 *
2668 * Input:
2669 *     ha:     adapter state pointer.
2670 *     sp:     SRB structure pointer.
2671 *
2672 * Returns:
2673 *     ql local function return status code.
2674 *
2675 * Context:
2676 *     Interrupt or Kernel context, no mailbox commands allowed.
2677 */
2678 static int
2679 ql_abort_cmd_iocb(ql_adapter_state_t *ha, ql_srb_t *sp)
2680 {
2681     ql_mbx_iocb_t    *pkt;
2682     int              rval;
2683     uint32_t         pkt_size;
2684     uint16_t         comp_status;
2685     ql_tgt_t         *tq = sp->lun_queue->target_queue;
2686
2687     QL_PRINT_3(CE_CONT, "(%d): started\n", ha->instance);
2688
2689     pkt_size = sizeof(ql_mbx_iocb_t);
2690     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
2691     if ((pkt = kmem_zalloc(pkt_size, KM_SLEEP)) == NULL) {
2692         EL(ha, "failed, kmem_zalloc\n");
2693         return (QL_MEMORY_ALLOC_FAILED);
2694     }
2695
2696     pkt->abo.entry_type = ABORT_CMD_TYPE;
2697     pkt->abo.entry_count = 1;
2698     pkt->abo.n_port_hdl = (uint16_t)LE_16(tq->loop_id);
2699     if (!CFG_IST(ha, CFG_CTRL_8021)) {
2700         pkt->abo.options = AF_NO_ABTS;
2701     }
2702     pkt->abo.cmd_handle = LE_32(sp->handle);
2703     pkt->abo.target_id[0] = tq->d_id.b.al_pa;
2704     pkt->abo.target_id[1] = tq->d_id.b.area;
2705     pkt->abo.target_id[2] = tq->d_id.b.domain;
2706     pkt->abo.vp_index = ha->vp_index;
2707
2708     rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt, pkt_size);
2709
2710     if (rval == QL_SUCCESS) {
2711         if ((pkt->abo.entry_status & 0x3c) != 0) {
2712             EL(ha, "failed, entry_status=%xh, d_id=%xh\n",
2713                pkt->abo.entry_status, tq->d_id.b24);
2714             rval = QL_FUNCTION_PARAMETER_ERROR;
2715         } else {
2716             comp_status = (uint16_t)LE_16(pkt->abo.n_port_hdl);
2717             if (comp_status != CS_COMPLETE) {
2718                 EL(ha, "failed, comp_status=%xh, d_id=%xh\n",
2719                    comp_status, tq->d_id.b24);
2720                 rval = QL_FUNCTION_FAILED;
2721             }
2722         }
2723     }
2724
2725     kmem_free(pkt, pkt_size);
2726
2727     if (rval != QL_SUCCESS) {
2728         EL(ha, "failed=%xh, d_id=%xh\n", rval, tq->d_id.b24);
2729     } else {
2730         /*EMPTY*/
2731     }

```

```

2727         QL_PRINT_3(CE_CONT, "(%d): done\n", ha->instance);
2728     }
2729
2730     return (rval);
2731 }

```

unchanged portion omitted

```

*****
50767 Thu Oct 23 11:04:55 2014
new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_nx.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
unchanged_portion_omitted

1370 static int
1371 ql_802l_pinit_from_rom(ql_adapter_state_t *ha)
1372 {
1373     int                init_delay = 0;
1374     struct crb_addr_pair *buf;
1375     uint32_t           offset, off, i, n, addr, val;

1377     /* Grab the lock so that no one can read flash when we reset the chip */
1378     (void) ql_802l_rom_lock(ha);
1379     ql_802l_wr_32(ha, UNM_ROMUSB_GLB_SW_RESET, 0xffffffff);
1380     /* Just in case it was held when we reset the chip */
1381     ql_802l_rom_unlock(ha);

1383     if (ql_802l_rom_fast_read(ha, 0, &n) != 0 || n != 0xcafecafe ||
1384         ql_802l_rom_fast_read(ha, 4, &n) != 0) {
1385         EL(ha, "ERROR Reading crb_init area: n: %08x\n", n);
1386         return (-1);
1387     }
1388     offset = n & 0xffff;
1389     n = (n >> 16) & 0xffff;
1390     if (n >= 1024) {
1391         EL(ha, "n=0x%x Error! NetXen card flash not initialized\n", n);
1392         return (-1);
1393     }

1395     buf = kmem_zalloc(n * sizeof (struct crb_addr_pair), KM_SLEEP);
1396     if (buf == NULL) {
1397         EL(ha, "Unable to zalloc memory\n");
1398         return (-1);
1399     }

1397     for (i = 0; i < n; i++) {
1398         if (ql_802l_rom_fast_read(ha, 8 * i + 4 * offset, &val) != 0 ||
1399             ql_802l_rom_fast_read(ha, 8 * i + 4 * offset + 4, &addr) !=
1400             0) {
1401             kmem_free(buf, n * sizeof (struct crb_addr_pair));
1402             EL(ha, "ql_802l_rom_fast_read != 0 to zalloc memory\n");
1403             return (-1);
1404         }

1406         buf[i].addr = addr;
1407         buf[i].data = val;
1408     }

1410     for (i = 0; i < n; i++) {
1411         off = ql_802l_decode_crb_addr(ha, buf[i].addr);
1412         if (off == ADDR_ERROR) {
1413             EL(ha, "Err: Unknown addr: 0x%lx\n", buf[i].addr);
1414             continue;
1415         }
1416         off += UNM_PCI_CRBSPACE;

1418         if (off & 1) {
1419             continue;
1420         }

1422         /* skipping cold reboot MAGIC */
1423         if (off == UNM_RAM_COLD_BOOT) {

```

```

1424             continue;
1425         }
1426         if (off == (UNM_CRB_I2C0 + 0x1c)) {
1427             continue;
1428         }
1429         /* do not reset PCI */
1430         if (off == (ROMUSB_GLB + 0xbc)) {
1431             continue;
1432         }
1433         if (off == (ROMUSB_GLB + 0xa8)) {
1434             continue;
1435         }
1436         if (off == (ROMUSB_GLB + 0xc8)) { /* core clock */
1437             continue;
1438         }
1439         if (off == (ROMUSB_GLB + 0x24)) { /* MN clock */
1440             continue;
1441         }
1442         if (off == (ROMUSB_GLB + 0x1c)) { /* MS clock */
1443             continue;
1444         }
1445         if ((off & 0x0ff00000) == UNM_CRB_DDR_NET) {
1446             continue;
1447         }
1448         if (off == (UNM_CRB_PEG_NET_1 + 0x18) &&
1449             !NX_IS_REVISION_P3PLUS(ha->rev_id)) {
1450             buf[i].data = 0x1020;
1451         }
1452         /* skip the function enable register */
1453         if (off == UNM_PCIE_REG(PCIE_SETUP_FUNCTION)) {
1454             continue;
1455         }
1456         if (off == UNM_PCIE_REG(PCIE_SETUP_FUNCTION2)) {
1457             continue;
1458         }
1459         if ((off & 0x0ff00000) == UNM_CRB_SMB) {
1460             continue;
1461         }

1463         /* After writing this register, HW needs time for CRB */
1464         /* to quiet down (else crb_window returns 0xffffffff) */
1465         init_delay = 1;
1466         if (off == UNM_ROMUSB_GLB_SW_RESET) {
1467             init_delay = 100; /* Sleep 1000 msecs */
1468         }

1470         ql_802l_wr_32(ha, off, buf[i].data);

1472         delay(init_delay);
1473     }
1474     kmem_free(buf, n * sizeof (struct crb_addr_pair));

1476     /* disable_peg_cache_all */

1478     /* p2dn replyCount */
1479     ql_802l_wr_32(ha, UNM_CRB_PEG_NET_D + 0xec, 0x1e);
1480     /* disable_peg_cache 0 */
1481     ql_802l_wr_32(ha, UNM_CRB_PEG_NET_D + 0x4c, 8);
1482     /* disable_peg_cache 1 */
1483     ql_802l_wr_32(ha, UNM_CRB_PEG_NET_I + 0x4c, 8);

1485     /* peg_clr_all */
1486     /* peg_clr 0 */
1487     ql_802l_wr_32(ha, UNM_CRB_PEG_NET_0 + 0x8, 0);
1488     ql_802l_wr_32(ha, UNM_CRB_PEG_NET_0 + 0xc, 0);
1489     /* peg_clr 1 */

```

new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_nx.c

3

```
1490     ql_8021_wr_32(ha, UNM_CRB_PEG_NET_1 + 0x8, 0);
1491     ql_8021_wr_32(ha, UNM_CRB_PEG_NET_1 + 0xc, 0);
1492     /* peg_clr 2 */
1493     ql_8021_wr_32(ha, UNM_CRB_PEG_NET_2 + 0x8, 0);
1494     ql_8021_wr_32(ha, UNM_CRB_PEG_NET_2 + 0xc, 0);
1495     /* peg_clr 3 */
1496     ql_8021_wr_32(ha, UNM_CRB_PEG_NET_3 + 0x8, 0);
1497     ql_8021_wr_32(ha, UNM_CRB_PEG_NET_3 + 0xc, 0);
1499     return (0);
1500 }
_____unchanged_portion_omitted_____
```

```

*****
235829 Thu Oct 23 11:04:55 2014
new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_xioctl.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */

22 /* Copyright 2010 QLogic Corporation */

24 /*
25  * Copyright (c) 2008, 2010, Oracle and/or its affiliates. All rights reserved.
26 */

28 /*
29  * Copyright 2011 Nexenta Systems, Inc. All rights reserved.
30 */

32 #pragma ident      "Copyright 2010 QLogic Corporation; ql_xioctl.c"

34 /*
35  * ISP2xxx Solaris Fibre Channel Adapter (FCA) driver source file.
36  *
37  * *****
38  * *
39  * *                NOTICE
40  * *                COPYRIGHT (C) 1996-2010 QLOGIC CORPORATION
41  * *                ALL RIGHTS RESERVED
42  * *
43  * *****
44  *
45  */

47 #include <ql_apps.h>
48 #include <ql_api.h>
49 #include <ql_debug.h>
50 #include <ql_init.h>
51 #include <ql_iocb.h>
52 #include <ql_ioctl.h>
53 #include <ql_mbx.h>
54 #include <ql_xioctl.h>

56 /*
57  * Local data
58 */

60 /*

```

```

61  * Local prototypes
62  */
63 static int ql_sdm_ioctl(ql_adapter_state_t *, int, void *, int);
64 static int ql_sdm_setup(ql_adapter_state_t *, EXT_IOCTL **, void *, int,
65     boolean_t (*)(EXT_IOCTL *));
66 static boolean_t ql_validate_signature(EXT_IOCTL *);
67 static int ql_sdm_return(ql_adapter_state_t *, EXT_IOCTL *, void *, int);
68 static void ql_query(ql_adapter_state_t *, EXT_IOCTL *, int);
69 static void ql_qry_hba_node(ql_adapter_state_t *, EXT_IOCTL *, int);
70 static void ql_qry_hba_port(ql_adapter_state_t *, EXT_IOCTL *, int);
71 static void ql_qry_disc_port(ql_adapter_state_t *, EXT_IOCTL *, int);
72 static void ql_qry_disc_tgt(ql_adapter_state_t *, EXT_IOCTL *, int);
73 static void ql_qry_fw(ql_adapter_state_t *, EXT_IOCTL *, int);
74 static void ql_qry_chip(ql_adapter_state_t *, EXT_IOCTL *, int);
75 static void ql_qry_driver(ql_adapter_state_t *, EXT_IOCTL *, int);
76 static void ql_fct(ql_adapter_state_t *, EXT_IOCTL *, int);
77 static void ql_aen_reg(ql_adapter_state_t *, EXT_IOCTL *, int);
78 static void ql_aen_get(ql_adapter_state_t *, EXT_IOCTL *, int);
79 static void ql_scsi_passthru(ql_adapter_state_t *, EXT_IOCTL *, int);
80 static void ql_wwpn_to_scsiaddr(ql_adapter_state_t *, EXT_IOCTL *, int);
81 static void ql_host_idx(ql_adapter_state_t *, EXT_IOCTL *, int);
82 static void ql_host_drvname(ql_adapter_state_t *, EXT_IOCTL *, int);
83 static void ql_read_nvram(ql_adapter_state_t *, EXT_IOCTL *, int);
84 static void ql_write_nvram(ql_adapter_state_t *, EXT_IOCTL *, int);
85 static void ql_read_flash(ql_adapter_state_t *, EXT_IOCTL *, int);
86 static void ql_write_flash(ql_adapter_state_t *, EXT_IOCTL *, int);
87 static void ql_write_vpd(ql_adapter_state_t *, EXT_IOCTL *, int);
88 static void ql_read_vpd(ql_adapter_state_t *, EXT_IOCTL *, int);
89 static void ql_diagnostic_loopback(ql_adapter_state_t *, EXT_IOCTL *, int);
90 static void ql_send_els_rnid(ql_adapter_state_t *, EXT_IOCTL *, int);
91 static void ql_set_host_data(ql_adapter_state_t *, EXT_IOCTL *, int);
92 static void ql_get_host_data(ql_adapter_state_t *, EXT_IOCTL *, int);
93 static void ql_qry_cna_port(ql_adapter_state_t *, EXT_IOCTL *, int);

95 static int ql_lun_count(ql_adapter_state_t *, ql_tgt_t *);
96 static int ql_report_lun(ql_adapter_state_t *, ql_tgt_t *);
97 static int ql_inq_scan(ql_adapter_state_t *, ql_tgt_t *, int);
98 static int ql_inq(ql_adapter_state_t *, ql_tgt_t *, int, ql_mbx_ioch_t *,
99     uint8_t);
100 static uint32_t ql_get_buffer_data(caddr_t, caddr_t, uint32_t, int);
101 static uint32_t ql_send_buffer_data(caddr_t, caddr_t, uint32_t, int);
102 static int ql_24xx_flash_desc(ql_adapter_state_t *);
103 static int ql_setup_flash(ql_adapter_state_t *);
104 static ql_tgt_t *ql_find_port(ql_adapter_state_t *, uint8_t *, uint16_t);
105 static int ql_flash_fcode_load(ql_adapter_state_t *, void *, uint32_t, int);
106 static int ql_flash_fcode_dump(ql_adapter_state_t *, void *, uint32_t,
107     uint32_t, int);
108 static int ql_program_flash_address(ql_adapter_state_t *, uint32_t,
109     uint8_t);
110 static void ql_set_rnid_parameters(ql_adapter_state_t *, EXT_IOCTL *, int);
111 static void ql_get_rnid_parameters(ql_adapter_state_t *, EXT_IOCTL *, int);
112 static int ql_reset_statistics(ql_adapter_state_t *, EXT_IOCTL *);
113 static void ql_get_statistics(ql_adapter_state_t *, EXT_IOCTL *, int);
114 static void ql_get_statistics_fc(ql_adapter_state_t *, EXT_IOCTL *, int);
115 static void ql_get_statistics_fc4(ql_adapter_state_t *, EXT_IOCTL *, int);
116 static void ql_set_led_state(ql_adapter_state_t *, EXT_IOCTL *, int);
117 static void ql_get_led_state(ql_adapter_state_t *, EXT_IOCTL *, int);
118 static void ql_drive_led(ql_adapter_state_t *, uint32_t);
119 static uint32_t ql_setup_led(ql_adapter_state_t *);
120 static uint32_t ql_wrapup_led(ql_adapter_state_t *);
121 static void ql_get_port_summary(ql_adapter_state_t *, EXT_IOCTL *, int);
122 static void ql_get_target_id(ql_adapter_state_t *, EXT_IOCTL *, int);
123 static void ql_get_sfp(ql_adapter_state_t *, EXT_IOCTL *, int);
124 static int ql_dump_sfp(ql_adapter_state_t *, void *, int);
125 static ql_fcache_t *ql_setup_fnode(ql_adapter_state_t *);
126 static void ql_get_fcache(ql_adapter_state_t *, EXT_IOCTL *, int);

```

```

127 static void ql_get_fcache_ex(ql_adapter_state_t *, EXT_IOCTL *, int);
128 void ql_update_fcache(ql_adapter_state_t *, uint8_t *, uint32_t);
129 static int ql_check_pci(ql_adapter_state_t *, ql_fcache_t *, uint32_t *);
130 static void ql_flash_layout_table(ql_adapter_state_t *, uint32_t);
131 static void ql_processflt(ql_adapter_state_t *, uint32_t);
132 static void ql_flash_nvram_defaults(ql_adapter_state_t *);
133 static void ql_port_param(ql_adapter_state_t *, EXT_IOCTL *, int);
134 static int ql_check_pci(ql_adapter_state_t *, ql_fcache_t *, uint32_t *);
135 static void ql_get_pci_data(ql_adapter_state_t *, EXT_IOCTL *, int);
136 static void ql_get_fwfcetrace(ql_adapter_state_t *, EXT_IOCTL *, int);
137 static void ql_get_fwexttrace(ql_adapter_state_t *, EXT_IOCTL *, int);
138 static void ql_menlo_reset(ql_adapter_state_t *, EXT_IOCTL *, int);
139 static void ql_menlo_get_fw_version(ql_adapter_state_t *, EXT_IOCTL *, int);
140 static void ql_menlo_update_fw(ql_adapter_state_t *, EXT_IOCTL *, int);
141 static void ql_menlo_manage_info(ql_adapter_state_t *, EXT_IOCTL *, int);
142 static int ql_suspend_hba(ql_adapter_state_t *, uint32_t);
143 static void ql_restart_hba(ql_adapter_state_t *);
144 static void ql_get_vp_cnt_id(ql_adapter_state_t *, EXT_IOCTL *, int);
145 static void ql_vp_ioctl(ql_adapter_state_t *, EXT_IOCTL *, int);
146 static void ql_gry_vport(ql_adapter_state_t *, EXT_IOCTL *, int);
147 static void ql_access_flash(ql_adapter_state_t *, EXT_IOCTL *, int);
148 static void ql_reset_cmd(ql_adapter_state_t *, EXT_IOCTL *);
149 static void ql_update_flash_caches(ql_adapter_state_t *);
150 static void ql_get_dcbx_parameters(ql_adapter_state_t *, EXT_IOCTL *, int);
151 static void ql_get_xgmac_statistics(ql_adapter_state_t *, EXT_IOCTL *, int);
152 static void ql_get_fcf_list(ql_adapter_state_t *, EXT_IOCTL *, int);
153 static void ql_get_resource_counts(ql_adapter_state_t *, EXT_IOCTL *, int);
154 static void ql_gry_adapter_versions(ql_adapter_state_t *, EXT_IOCTL *, int);
155 static int ql_set_loop_point(ql_adapter_state_t *, uint16_t);

157 /* ***** */
158 /* External IOCTL support. */
159 /* ***** */

161 /*
162 * ql_alloc_xioctl_resource
163 * Allocates resources needed by module code.
164 *
165 * Input:
166 * ha: adapter state pointer.
167 *
168 * Returns:
169 * SYS_ERRNO
170 *
171 * Context:
172 * Kernel context.
173 */
174 int
175 ql_alloc_xioctl_resource(ql_adapter_state_t *ha)
176 {
177     ql_xioctl_t *xp;

179     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

181     if (ha->xioctl != NULL) {
182         QL_PRINT_9(CE_CONT, "(%d): already allocated done\n",
183             ha->instance);
184         return (0);
185     }

187     xp = kmem_zalloc(sizeof(ql_xioctl_t), KM_SLEEP);
188     if (xp == NULL) {
189         EL(ha, "failed, kmem_zalloc\n");
190         return (ENOMEM);
191     }
192     ha->xioctl = xp;

```

```

190     /* Allocate AEN tracking buffer */
191     xp->aen_tracking_queue = kmem_zalloc(EXT_DEF_MAX_AEN_QUEUE *
192         sizeof(EXT_ASYNC_EVENT), KM_SLEEP);
193     if (xp->aen_tracking_queue == NULL) {
194         EL(ha, "failed, kmem_zalloc-2\n");
195         ql_free_xioctl_resource(ha);
196         return (ENOMEM);
197     }

199     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
200     return (0);
201 }

203 /*
204 * ql_fcct
205 * IOCTL management server FC-CT passthrough.
206 *
207 * Input:
208 * ha: adapter state pointer.
209 * cmd: User space CT arguments pointer.
210 * mode: flags.
211 *
212 * Returns:
213 * None, request status indicated in cmd->Status.
214 *
215 * Context:
216 * Kernel context.
217 */
218 static void
219 ql_fcct(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
220 {
221     ql_mbx_iocb_t *pkt;
222     ql_mbx_data_t nr;
223     dma_mem_t *dma_mem;
224     pld_t pld;
225     uint32_t pkt_size, pld_byte_cnt, *long_ptr;
226     int rval;
227     ql_ct_iu_preamble_t *ct;
228     ql_xioctl_t *xp = ha->xioctl;
229     tq_t tq;
230     uint16_t comp_status, loop_id;

232     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

234     /* Get CT argument structure. */
235     if ((ha->topology & QL_SNS_CONNECTION) == 0) {
236         EL(ha, "failed, No switch\n");
237         cmd->Status = EXT_STATUS_DEV_NOT_FOUND;
238         cmd->ResponseLen = 0;
239         return;
240     }

242     if (DRIVER_SUSPENDED(ha)) {
243         EL(ha, "failed, LOOP_NOT_READY\n");
244         cmd->Status = EXT_STATUS_BUSY;
245         cmd->ResponseLen = 0;
246         return;
247     }

249     /* Login management server device. */
250     if ((xp->flags & QL_MGMT_SERVER_LOGIN) == 0) {
251         tq.d_id.b.al_pa = 0xf;
252         tq.d_id.b.area = 0xff;

```



```

1488         tq.d_id.b.domain = 0xff;
1489         tq.loop_id = (uint16_t)(CFG_IST(ha, CFG_CTRL_24258081) ?
1490             MANAGEMENT_SERVER_24XX_LOOP_ID :
1491             MANAGEMENT_SERVER_LOOP_ID);
1492         rval = ql_login_fport(ha, &tq, tq.loop_id, LFF_NO_PRLI, &mr);
1493         if (rval != QL_SUCCESS) {
1494             EL(ha, "failed, server login\n");
1495             cmd->Status = EXT_STATUS_DEV_NOT_FOUND;
1496             cmd->ResponseLen = 0;
1497             return;
1498         } else {
1499             xp->flags |= QL_MGMT_SERVER_LOGIN;
1500         }
1501     }

1503     QL_PRINT_9(CE_CONT, "%d): cmd\n", ha->instance);
1504     QL_DUMP_9(cmd, 8, sizeof (EXT_IOCTL));

1506     /* Allocate a DMA Memory Descriptor */
1507     dma_mem = (dma_mem_t *)kmem_zalloc(sizeof (dma_mem_t), KM_SLEEP);
1508     if (dma_mem == NULL) {
1509         EL(ha, "failed, kmem_zalloc\n");
1510         cmd->Status = EXT_STATUS_NO_MEMORY;
1511         cmd->ResponseLen = 0;
1512         return;
1513     }
1514     /* Determine maximum buffer size. */
1515     if (cmd->RequestLen < cmd->ResponseLen) {
1516         pld_byte_cnt = cmd->ResponseLen;
1517     } else {
1518         pld_byte_cnt = cmd->RequestLen;
1519     }

1521     /* Allocate command block. */
1522     pkt_size = (uint32_t)(sizeof (ql_mbx_iocb_t) + pld_byte_cnt);
1523     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
1524     if (pkt == NULL) {
1525         EL(ha, "failed, kmem_zalloc\n");
1526         cmd->Status = EXT_STATUS_NO_MEMORY;
1527         cmd->ResponseLen = 0;
1528         return;
1529     }
1530     pld = (caddr_t)pkt + sizeof (ql_mbx_iocb_t);

1532     /* Get command payload data. */
1533     if (ql_get_buffer_data((caddr_t)(uintptr_t)cmd->RequestAdr, pld,
1534         cmd->RequestLen, mode) != cmd->RequestLen) {
1535         EL(ha, "failed, get_buffer_data\n");
1536         kmem_free(pkt, pkt_size);
1537         cmd->Status = EXT_STATUS_COPY_ERR;
1538         cmd->ResponseLen = 0;
1539         return;
1540     }

1542     /* Get DMA memory for the IOCB */
1543     if (ql_get_dma_mem(ha, dma_mem, pkt_size, LITTLE_ENDIAN_DMA,
1544         QL_DMA_RING_ALIGN) != QL_SUCCESS) {
1545         cmn_err(CE_WARN, "%s(%d): DMA memory "
1546             "alloc failed", QL_NAME, ha->instance);
1547         kmem_free(pkt, pkt_size);
1548         kmem_free(dma_mem, sizeof (dma_mem_t));
1549         cmd->Status = EXT_STATUS_MS_NO_RESPONSE;
1550         cmd->ResponseLen = 0;
1551         return;
1552     }

```

```

1548         /* Copy out going payload data to IOCB DMA buffer. */
1549         ddi_rep_put8(dma_mem->acc_handle, (uint8_t *)pld,
1550             (uint8_t *)dma_mem->bp, pld_byte_cnt, DDI_DEV_AUTOINCR);

1552         /* Sync IOCB DMA buffer. */
1553         (void) ddi_dma_sync(dma_mem->dma_handle, 0, pld_byte_cnt,
1554             DDI_DMA_SYNC_FORDEV);

1556         /*
1557          * Setup IOCB
1558          */
1559         ct = (ql_ct_iu_preamble_t *)pld;
1560         if (CFG_IST(ha, CFG_CTRL_24258081)) {
1561             pkt->ms24.entry_type = CT_PASSTHRU_TYPE;
1562             pkt->ms24.entry_count = 1;

1564             pkt->ms24.vp_index = ha->vp_index;

1566             /* Set loop ID */
1567             pkt->ms24.n_port_hdl = (uint16_t)
1568                 (ct->gs_type == GS_TYPE_DIR_SERVER ?
1569                 LE_16(SNS_24XX_HDL) :
1570                 LE_16(MANAGEMENT_SERVER_24XX_LOOP_ID));

1572             /* Set ISP command timeout. */
1573             pkt->ms24.timeout = LE_16(120);

1575             /* Set cmd/response data segment counts. */
1576             pkt->ms24.cmd_dseg_count = LE_16(1);
1577             pkt->ms24.resp_dseg_count = LE_16(1);

1579             /* Load ct cmd byte count. */
1580             pkt->ms24.cmd_byte_count = LE_32(cmd->RequestLen);

1582             /* Load ct rsp byte count. */
1583             pkt->ms24.resp_byte_count = LE_32(cmd->ResponseLen);

1585             long_ptr = (uint32_t *)&pkt->ms24.dseg_0_address;

1587             /* Load MS command entry data segments. */
1588             *long_ptr++ = (uint32_t)
1589                 LE_32(LSD(dma_mem->cookie.dmac_laddress));
1590             *long_ptr++ = (uint32_t)
1591                 LE_32(MSD(dma_mem->cookie.dmac_laddress));
1592             *long_ptr++ = (uint32_t)(LE_32(cmd->RequestLen));

1594             /* Load MS response entry data segments. */
1595             *long_ptr++ = (uint32_t)
1596                 LE_32(LSD(dma_mem->cookie.dmac_laddress));
1597             *long_ptr++ = (uint32_t)
1598                 LE_32(MSD(dma_mem->cookie.dmac_laddress));
1599             *long_ptr = (uint32_t)LE_32(cmd->ResponseLen);

1601             rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt,
1602                 sizeof (ql_mbx_iocb_t));

1604             comp_status = (uint16_t)LE_16(pkt->sts24.comp_status);
1605             if (comp_status == CS_DATA_UNDEERRUN) {
1606                 if ((BE_16(ct->max_residual_size) == 0) {
1607                     comp_status = CS_COMPLETE;
1608                 }
1609             }

1611             if (rval != QL_SUCCESS || (pkt->sts24.entry_status & 0x3c) !=
1612                 0) {
1613                 EL(ha, "failed, I/O timeout or "

```

```

1614         "es=%xh, ss_l=%xh, rval=%xh\n",
1615         pkt->sts24.entry_status,
1616         pkt->sts24.scsi_status_l, rval);
1617     kmem_free(pkt, pkt_size);
1618     ql_free_dma_resource(ha, dma_mem);
1619     kmem_free(dma_mem, sizeof (dma_mem_t));
1620     cmd->Status = EXT_STATUS_MS_NO_RESPONSE;
1621     cmd->ResponseLen = 0;
1622     return;
1623 }
1624 } else {
1625     pkt->ms.entry_type = MS_TYPE;
1626     pkt->ms.entry_count = 1;

1628     /* Set loop ID */
1629     loop_id = (uint16_t)(ct->gs_type == GS_TYPE_DIR_SERVER ?
1630     SIMPLE_NAME_SERVER_LOOP_ID : MANAGEMENT_SERVER_LOOP_ID);
1631     if (CFG_IST(ha, CFG_EXT_FW_INTERFACE)) {
1632         pkt->ms.loop_id_l = LSB(loop_id);
1633         pkt->ms.loop_id_h = MSB(loop_id);
1634     } else {
1635         pkt->ms.loop_id_h = LSB(loop_id);
1636     }

1638     /* Set ISP command timeout. */
1639     pkt->ms.timeout = LE_16(120);

1641     /* Set data segment counts. */
1642     pkt->ms.cmd_dseg_count_l = 1;
1643     pkt->ms.total_dseg_count = LE_16(2);

1645     /* Response total byte count. */
1646     pkt->ms.resp_byte_count = LE_32(cmd->ResponseLen);
1647     pkt->ms.dseg_l_length = LE_32(cmd->ResponseLen);

1649     /* Command total byte count. */
1650     pkt->ms.cmd_byte_count = LE_32(cmd->RequestLen);
1651     pkt->ms.dseg_0_length = LE_32(cmd->RequestLen);

1653     /* Load command/response data segments. */
1654     pkt->ms.dseg_0_address[0] = (uint32_t)
1655     LE_32(LSD(dma_mem->cookie.dmac_laddress));
1656     pkt->ms.dseg_0_address[1] = (uint32_t)
1657     LE_32(MSD(dma_mem->cookie.dmac_laddress));
1658     pkt->ms.dseg_l_address[0] = (uint32_t)
1659     LE_32(LSD(dma_mem->cookie.dmac_laddress));
1660     pkt->ms.dseg_l_address[1] = (uint32_t)
1661     LE_32(MSD(dma_mem->cookie.dmac_laddress));

1663     rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt,
1664     sizeof (ql_mbx_iocb_t));

1666     comp_status = (uint16_t)LE_16(pkt->sts.comp_status);
1667     if (comp_status == CS_DATA_UNDERRUN) {
1668         if ((BE_16(ct->max_residual_size) == 0) {
1669             comp_status = CS_COMPLETE;
1670         }
1671     }
1672     if (rval != QL_SUCCESS || (pkt->sts.entry_status & 0x7e) != 0) {
1673         EL(ha, "failed, I/O timeout or "
1674         "es=%xh, rval=%xh\n", pkt->sts.entry_status, rval);
1675         kmem_free(pkt, pkt_size);
1676         ql_free_dma_resource(ha, dma_mem);
1677         kmem_free(dma_mem, sizeof (dma_mem_t));
1678         cmd->Status = EXT_STATUS_MS_NO_RESPONSE;
1679         cmd->ResponseLen = 0;

```

```

1680         return;
1681     }
1682 }

1684 /* Sync in coming DMA buffer. */
1685 (void) ddi_dma_sync(dma_mem->dma_handle, 0,
1686     pld_byte_cnt, DDI_DMA_SYNC_FORKERNEL);
1687 /* Copy in coming DMA data. */
1688 ddi_rep_get8(dma_mem->acc_handle, (uint8_t *)pld,
1689     (uint8_t *)dma_mem->bp, pld_byte_cnt,
1690     DDI_DEV_AUTOINCR);

1692 /* Copy response payload from DMA buffer to application. */
1693 if (cmd->ResponseLen != 0) {
1694     QL_PRINT_9(CE_CONT, "(&d): ResponseLen=%d\n", ha->instance,
1695     cmd->ResponseLen);
1696     QL_DUMP_9(pld, 8, cmd->ResponseLen);

1698     /* Send response payload. */
1699     if (ql_send_buffer_data(pld,
1700     (caddr_t)(uintptr_t)cmd->ResponseAdr,
1701     cmd->ResponseLen, mode) != cmd->ResponseLen) {
1702         EL(ha, "failed, send_buffer_data\n");
1703         cmd->Status = EXT_STATUS_COPY_ERR;
1704         cmd->ResponseLen = 0;
1705     }
1706 }

1708     kmem_free(pkt, pkt_size);
1709     ql_free_dma_resource(ha, dma_mem);
1710     kmem_free(dma_mem, sizeof (dma_mem_t));

1712     QL_PRINT_9(CE_CONT, "(&d): done\n", ha->instance);
1713 }

    unchanged portion omitted

1940 /*
1941  * ql_scsi_passthru
1942  * IOCTL SCSI passthrough.
1943  *
1944  * Input:
1945  *     ha: adapter state pointer.
1946  *     cmd: User space SCSI command pointer.
1947  *     mode: flags.
1948  *
1949  * Returns:
1950  *     None, request status indicated in cmd->Status.
1951  *
1952  * Context:
1953  *     Kernel context.
1954  */
1955 static void
1956 ql_scsi_passthru(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
1957 {
1958     ql_mbx_iocb_t     *pkt;
1959     ql_mbx_data_t     mr;
1960     dma_mem_t         *dma_mem;
1961     caddr_t           pld;
1962     uint32_t          pkt_size, pld_size;
1963     uint16_t          qlnt, retries, cnt, cnt2;
1964     uint8_t           *name;
1965     EXT_FC_SCSI_PASSTHRU *ufc_req;
1966     EXT_SCSI_PASSTHRU *usp_req;
1967     int               rval;
1968     union _passthru {
1969         EXT_SCSI_PASSTHRU     sp_cmd;

```

```

1970         EXT_FC_SCSI_PASSTHRU   fc_cmd;
1971     } pt_req; /* Passthru request */
1972     uint32_t      status, sense_sz = 0;
1973     ql_tgt_t      *tq = NULL;
1974     EXT_SCSI_PASSTHRU *sp_req = &pt_req.sp_cmd;
1975     EXT_FC_SCSI_PASSTHRU *fc_req = &pt_req.fc_cmd;

1977 /* SCSI request struct for SCSI passthrough IOs. */
1978 struct {
1979     uint16_t      lun;
1980     uint16_t      sense_length; /* Sense buffer size */
1981     size_t        resid; /* Residual */
1982     uint8_t       *cdbp; /* Requestor's CDB */
1983     uint8_t       *u_sense; /* Requestor's sense buffer */
1984     uint8_t       cdb_len; /* Requestor's CDB length */
1985     uint8_t       direction;
1986 } scsi_req;

1988 struct {
1989     uint8_t        *rsp_info;
1990     uint8_t        *req_sense_data;
1991     uint32_t       residual_length;
1992     uint32_t       rsp_info_length;
1993     uint32_t       req_sense_length;
1994     uint16_t       comp_status;
1995     uint8_t        state_flags_l;
1996     uint8_t        state_flags_h;
1997     uint8_t        scsi_status_l;
1998     uint8_t        scsi_status_h;
1999 } sts;

2001 QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

2003 /* Verify Sub Code and set cnt to needed request size. */
2004 if (cmd->SubCode == EXT_SC_SEND_SCSI_PASSTHRU) {
2005     pld_size = sizeof(EXT_SCSI_PASSTHRU);
2006 } else if (cmd->SubCode == EXT_SC_SEND_FC_SCSI_PASSTHRU) {
2007     pld_size = sizeof(EXT_FC_SCSI_PASSTHRU);
2008 } else {
2009     EL(ha, "failed, invalid SubCode=%xh\n", cmd->SubCode);
2010     cmd->Status = EXT_STATUS_UNSUPPORTED_SUBCODE;
2011     cmd->ResponseLen = 0;
2012     return;
2013 }

2015 dma_mem = (dma_mem_t *)kmem_zalloc(sizeof(dma_mem_t), KM_SLEEP);
2016 if (dma_mem == NULL) {
2017     EL(ha, "failed, kmem_zalloc\n");
2018     cmd->Status = EXT_STATUS_NO_MEMORY;
2019     cmd->ResponseLen = 0;
2020     return;
2021 }
2022 /* Verify the size of and copy in the passthru request structure. */
2023 if (cmd->RequestLen != pld_size) {
2024     /* Return error */
2025     EL(ha, "failed, RequestLen != cnt, is=%xh, expected=%xh\n",
2026         cmd->RequestLen, pld_size);
2027     cmd->Status = EXT_STATUS_INVALID_PARAM;
2028     cmd->DetailStatus = EXT_DSTATUS_REQUEST_LEN;
2029     cmd->ResponseLen = 0;
2030     return;
2031 }

2033 if (ddi_copyin((void *) (uintptr_t) cmd->RequestAdr, &pt_req,
2034     pld_size, mode) != 0) {
2035     EL(ha, "failed, ddi_copyin\n");

```

```

2036         cmd->Status = EXT_STATUS_COPY_ERR;
2037         cmd->ResponseLen = 0;
2038         return;
2039     }

2041 /*
2042  * Find fc_port from SCSI PASSTHRU structure fill in the scsi_req
2043  * request data structure.
2044  */
2045 if (cmd->SubCode == EXT_SC_SEND_SCSI_PASSTHRU) {
2046     scsi_req.lun = sp_req->TargetAddr.Lun;
2047     scsi_req.sense_length = sizeof(sp_req->SenseData);
2048     scsi_req.cdbp = &sp_req->Cdb[0];
2049     scsi_req.cdb_len = sp_req->CdbLength;
2050     scsi_req.direction = sp_req->Direction;
2051     usp_req = (EXT_SCSI_PASSTHRU *) (uintptr_t) cmd->RequestAdr;
2052     scsi_req.u_sense = &usp_req->SenseData[0];
2053     cmd->DetailStatus = EXT_DSTATUS_TARGET;

2055     qlnt = QLNT_PORT;
2056     name = (uint8_t *) &sp_req->TargetAddr.Target;
2057     QL_PRINT_9(CE_CONT, "(%d): SubCode=%xh, Target=%lld\n",
2058         ha->instance, cmd->SubCode, sp_req->TargetAddr.Target);
2059     tq = ql_find_port(ha, name, qlnt);
2060 } else {
2061     /*
2062      * Must be FC PASSTHRU, verified above.
2063      */
2064     if (fc_req->FCScsiAddr.DestType == EXT_DEF_DESTTYPE_WWPN) {
2065         qlnt = QLNT_PORT;
2066         name = &fc_req->FCScsiAddr.DestAddr.WWPN[0];
2067         QL_PRINT_9(CE_CONT, "(%d): SubCode=%xh, "
2068             "wwpn=%02x%02x%02x%02x%02x%02x%02x%02x\n",
2069             ha->instance, cmd->SubCode, name[0], name[1],
2070             name[2], name[3], name[4], name[5], name[6],
2071             name[7]);
2072         tq = ql_find_port(ha, name, qlnt);
2073     } else if (fc_req->FCScsiAddr.DestType ==
2074         EXT_DEF_DESTTYPE_WWNN) {
2075         qlnt = QLNT_NODE;
2076         name = &fc_req->FCScsiAddr.DestAddr.WWNN[0];
2077         QL_PRINT_9(CE_CONT, "(%d): SubCode=%xh, "
2078             "wwnn=%02x%02x%02x%02x%02x%02x%02x%02x\n",
2079             ha->instance, cmd->SubCode, name[0], name[1],
2080             name[2], name[3], name[4], name[5], name[6],
2081             name[7]);
2082         tq = ql_find_port(ha, name, qlnt);
2083     } else if (fc_req->FCScsiAddr.DestType ==
2084         EXT_DEF_DESTTYPE_PORTID) {
2085         qlnt = QLNT_PID;
2086         name = &fc_req->FCScsiAddr.DestAddr.Id[0];
2087         QL_PRINT_9(CE_CONT, "(%d): SubCode=%xh, PID="
2088             "%02x%02x%02x\n", ha->instance, cmd->SubCode,
2089             name[0], name[1], name[2]);
2090         tq = ql_find_port(ha, name, qlnt);
2091     } else {
2092         EL(ha, "failed, SubCode=%xh invalid DestType=%xh\n",
2093             cmd->SubCode, fc_req->FCScsiAddr.DestType);
2094         cmd->Status = EXT_STATUS_INVALID_PARAM;
2095         cmd->ResponseLen = 0;
2096         return;
2097     }
2098     scsi_req.lun = fc_req->FCScsiAddr.Lun;
2099     scsi_req.sense_length = sizeof(fc_req->SenseData);
2100     scsi_req.cdbp = &sp_req->Cdb[0];
2101     scsi_req.cdb_len = sp_req->CdbLength;

```

```

2102         ufc_req = (EXT_FC_SCSI_PASSTHRU *) (uintptr_t) cmd->RequestAdr;
2103         scsi_req.u_sense = &ufc_req->SenseData[0];
2104         scsi_req.direction = fc_req->Direction;
2105     }

2107     if (tq == NULL || !VALID_TARGET_ID(ha, tq->loop_id)) {
2108         EL(ha, "failed, fc_port not found\n");
2109         cmd->Status = EXT_STATUS_DEV_NOT_FOUND;
2110         cmd->ResponseLen = 0;
2111         return;
2112     }

2114     if (tq->flags & TQF_NEED_AUTHENTICATION) {
2115         EL(ha, "target not available; loopid=%xh\n", tq->loop_id);
2116         cmd->Status = EXT_STATUS_DEVICE_OFFLINE;
2117         cmd->ResponseLen = 0;
2118         return;
2119     }

2121     /* Allocate command block. */
2122     if ((scsi_req.direction == EXT_DEF_SCSI_PASSTHRU_DATA_IN ||
2123         scsi_req.direction == EXT_DEF_SCSI_PASSTHRU_DATA_OUT) &&
2124         cmd->ResponseLen) {
2125         pld_size = cmd->ResponseLen;
2126         pkt_size = (uint32_t)(sizeof(ql_mbx_iocb_t) + pld_size);
2127         pkt = kmem_zalloc(pkt_size, KM_SLEEP);
2128         if (pkt == NULL) {
2129             EL(ha, "failed, kmem_zalloc\n");
2130             cmd->Status = EXT_STATUS_NO_MEMORY;
2131             cmd->ResponseLen = 0;
2132             return;
2133         }
2134         pld = (caddr_t)pkt + sizeof(ql_mbx_iocb_t);

2136     /* Get DMA memory for the IOCB */
2137     if (ql_get_dma_mem(ha, dma_mem, pld_size, LITTLE_ENDIAN_DMA,
2138         QL_DMA_DATA_ALIGN) != QL_SUCCESS) {
2139         cmn_err(CE_WARN, "%s(%d): request queue DMA memory "
2140             "alloc failed", QL_NAME, ha->instance);
2141         kmem_free(pkt, pkt_size);
2142         cmd->Status = EXT_STATUS_MS_NO_RESPONSE;
2143         cmd->ResponseLen = 0;
2144         return;
2145     }

2146     if (scsi_req.direction == EXT_DEF_SCSI_PASSTHRU_DATA_IN) {
2147         scsi_req.direction = (uint8_t)
2148             (CFG_IST(ha, CFG_CTRL_24258081) ?
2149              CF_RD : CF_DATA_IN | CF_STAG);
2150     } else {
2151         scsi_req.direction = (uint8_t)
2152             (CFG_IST(ha, CFG_CTRL_24258081) ?
2153              CF_WR : CF_DATA_OUT | CF_STAG);
2154         cmd->ResponseLen = 0;

2156     /* Get command payload. */
2157     if (ql_get_buffer_data(
2158         (caddr_t)(uintptr_t)cmd->ResponseAdr,
2159         pld, pld_size, mode) != pld_size) {
2160         EL(ha, "failed, get_buffer_data\n");
2161         cmd->Status = EXT_STATUS_COPY_ERR;

2162         kmem_free(pkt, pkt_size);
2163         ql_free_dma_resource(ha, dma_mem);
2164         kmem_free(dma_mem, sizeof(dma_mem_t));
2165         return;

```

```

2162     }

2164     /* Copy out going data to DMA buffer. */
2165     ddi_rep_put8(dma_mem->acc_handle, (uint8_t *)pld,
2166         (uint8_t *)dma_mem->bp, pld_size,
2167         DDI_DEV_AUTOINCR);

2169     /* Sync DMA buffer. */
2170     (void) ddi_dma_sync(dma_mem->dma_handle, 0,
2171         dma_mem->size, DDI_DMA_SYNC_FORDEV);
2172 }
2173 } else {
2174     scsi_req.direction = (uint8_t)
2175         (CFG_IST(ha, CFG_CTRL_24258081) ? 0 : CF_STAG);
2176     cmd->ResponseLen = 0;

2178     pkt_size = sizeof(ql_mbx_iocb_t);
2179     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
2180     if (pkt == NULL) {
2181         EL(ha, "failed, kmem_zalloc-2\n");
2182         cmd->Status = EXT_STATUS_NO_MEMORY;
2183         return;
2184     }
2185     pld = NULL;
2186     pld_size = 0;
2187 }

2188 /* retries = ha->port_down_retry_count; */
2189 retries = 1;
2190 cmd->Status = EXT_STATUS_OK;
2191 cmd->DetailStatus = EXT_DSTATUS_NOADNL_INFO;

2192 QL_PRINT_9(CE_CONT, "%d): SCSI cdb\n", ha->instance);
2193 QL_DUMP_9(scsi_req.cdbp, 8, scsi_req.cdb_len);

2194 do {
2195     if (DRIVER_SUSPENDED(ha)) {
2196         sts.comp_status = CS_LOOP_DOWN_ABORT;
2197         break;
2198     }

2199     if (CFG_IST(ha, CFG_CTRL_24258081)) {
2200         pkt->cmd24.entry_type = IOCB_CMD_TYPE_7;
2201         pkt->cmd24.entry_count = 1;

2202         /* Set LUN number */
2203         pkt->cmd24.fcp_lun[2] = LSB(scsi_req.lun);
2204         pkt->cmd24.fcp_lun[3] = MSB(scsi_req.lun);

2206         /* Set N_port handle */
2207         pkt->cmd24.n_port_hdl = (uint16_t)LE_16(tq->loop_id);

2209         /* Set VP Index */
2210         pkt->cmd24.vp_index = ha->vp_index;

2212         /* Set target ID */
2213         pkt->cmd24.target_id[0] = tq->d_id.b.al_pa;
2214         pkt->cmd24.target_id[1] = tq->d_id.b.area;
2215         pkt->cmd24.target_id[2] = tq->d_id.b.domain;

2217         /* Set ISP command timeout. */
2218         pkt->cmd24.timeout = (uint16_t)LE_16(15);

2220         /* Load SCSI CDB */
2221         ddi_rep_put8(ha->hba_buf.acc_handle, scsi_req.cdbp,
2222             pkt->cmd24.scsi_cdb, scsi_req.cdb_len,

```

```

2223         DDI_DEV_AUTOINCR);
2224     for (cnt = 0; cnt < MAX_CMDSZ;
2225         cnt = (uint16_t)(cnt + 4)) {
2226         ql_chg_endian((uint8_t *) &pkt->cmd24.scsi_cdb
2227             + cnt, 4);
2228     }
2229
2230     /* Set tag queue control flags */
2231     pkt->cmd24.task = TA_STAG;
2232
2233     if (pld_size) {
2234         /* Set transfer direction. */
2235         pkt->cmd24.control_flags = scsi_req.direction;
2236
2237         /* Set data segment count. */
2238         pkt->cmd24.dseg_count = LE_16(1);
2239
2240         /* Load total byte count. */
2241         pkt->cmd24.total_byte_count = LE_32(pld_size);
2242
2243         /* Load data descriptor. */
2244         pkt->cmd24.dseg_0_address[0] = (uint32_t)
2245             LE_32(LSD(dma_mem->cookie.dmac_laddress));
2246         pkt->cmd24.dseg_0_address[1] = (uint32_t)
2247             LE_32(MSD(dma_mem->cookie.dmac_laddress));
2248         pkt->cmd24.dseg_0_length = LE_32(pld_size);
2249     }
2250     } else if (CFG_IST(ha, CFG_ENABLE_64BIT_ADDRESSING)) {
2251         pkt->cmd3.entry_type = IOCB_CMD_TYPE_3;
2252         pkt->cmd3.entry_count = 1;
2253         if (CFG_IST(ha, CFG_EXT_FW_INTERFACE)) {
2254             pkt->cmd3.target_l = LSB(tq->loop_id);
2255             pkt->cmd3.target_h = MSB(tq->loop_id);
2256         } else {
2257             pkt->cmd3.target_h = LSB(tq->loop_id);
2258         }
2259         pkt->cmd3.lun_l = LSB(scsi_req.lun);
2260         pkt->cmd3.lun_h = MSB(scsi_req.lun);
2261         pkt->cmd3.control_flags_l = scsi_req.direction;
2262         pkt->cmd3.timeout = LE_16(15);
2263         for (cnt = 0; cnt < scsi_req.cdb_len; cnt++) {
2264             pkt->cmd3.scsi_cdb[cnt] = scsi_req.cdbp[cnt];
2265         }
2266         if (pld_size) {
2267             pkt->cmd3.dseg_count = LE_16(1);
2268             pkt->cmd3.byte_count = LE_32(pld_size);
2269             pkt->cmd3.dseg_0_address[0] = (uint32_t)
2270                 LE_32(LSD(dma_mem->cookie.dmac_laddress));
2271             pkt->cmd3.dseg_0_address[1] = (uint32_t)
2272                 LE_32(MSD(dma_mem->cookie.dmac_laddress));
2273             pkt->cmd3.dseg_0_length = LE_32(pld_size);
2274         }
2275     } else {
2276         pkt->cmd.entry_type = IOCB_CMD_TYPE_2;
2277         pkt->cmd.entry_count = 1;
2278         if (CFG_IST(ha, CFG_EXT_FW_INTERFACE)) {
2279             pkt->cmd.target_l = LSB(tq->loop_id);
2280             pkt->cmd.target_h = MSB(tq->loop_id);
2281         } else {
2282             pkt->cmd.target_h = LSB(tq->loop_id);
2283         }
2284         pkt->cmd.lun_l = LSB(scsi_req.lun);
2285         pkt->cmd.lun_h = MSB(scsi_req.lun);
2286         pkt->cmd.control_flags_l = scsi_req.direction;
2287         pkt->cmd.timeout = LE_16(15);
2288         for (cnt = 0; cnt < scsi_req.cdb_len; cnt++) {

```

```

2289         pkt->cmd.scsi_cdb[cnt] = scsi_req.cdbp[cnt];
2290     }
2291     if (pld_size) {
2292         pkt->cmd.dseg_count = LE_16(1);
2293         pkt->cmd.byte_count = LE_32(pld_size);
2294         pkt->cmd.dseg_0_address = (uint32_t)
2295             LE_32(LSD(dma_mem->cookie.dmac_laddress));
2296         pkt->cmd.dseg_0_length = LE_32(pld_size);
2297     }
2298 }
2299 /* Go issue command and wait for completion. */
2300 QL_PRINT_9(CE_CONT, "(%d): request pkt\n", ha->instance);
2301 QL_DUMP_9(pkt, 8, pkt_size);
2302
2303 status = ql_issue_mbx_iocb(ha, (caddr_t)pkt, pkt_size);
2304
2305 if (pld_size) {
2306     /* Sync in coming DMA buffer. */
2307     (void) ddi_dma_sync(dma_mem->dma_handle, 0,
2308         dma_mem->size, DDI_DMA_SYNC_FORKERNEL);
2309     /* Copy in coming DMA data. */
2310     ddi_rep_get8(dma_mem->acc_handle, (uint8_t *)pld,
2311         (uint8_t *)dma_mem->bp, pld_size,
2312         DDI_DEV_AUTOINCR);
2313 }
2314
2315 if (CFG_IST(ha, CFG_CTRL_24258081)) {
2316     pkt->sts24.entry_status = (uint8_t)
2317         (pkt->sts24.entry_status & 0x3c);
2318 } else {
2319     pkt->sts.entry_status = (uint8_t)
2320         (pkt->sts.entry_status & 0x7e);
2321 }
2322
2323 if (status == QL_SUCCESS && pkt->sts.entry_status != 0) {
2324     EL(ha, "failed, entry_status=%xh, d_id=%xh\n",
2325         pkt->sts.entry_status, tq->d_id.b24);
2326     status = QL_FUNCTION_PARAMETER_ERROR;
2327 }
2328
2329 sts.comp_status = (uint16_t)(CFG_IST(ha, CFG_CTRL_24258081) ?
2330     LE_16(pkt->sts24.comp_status) :
2331     LE_16(pkt->sts.comp_status));
2332
2333 /*
2334  * We have verified about all the request that can be so far.
2335  * Now we need to start verification of our ability to
2336  * actually issue the CDB.
2337  */
2338 if (DRIVER_SUSPENDED(ha)) {
2339     sts.comp_status = CS_LOOP_DOWN_ABORT;
2340     break;
2341 } else if (status == QL_SUCCESS &&
2342     (sts.comp_status == CS_PORT_LOGGED_OUT ||
2343     sts.comp_status == CS_PORT_UNAVAILABLE)) {
2344     EL(ha, "login retry d_id=%xh\n", tq->d_id.b24);
2345     if (tq->flags & TQF_FABRIC_DEVICE) {
2346         rval = ql_login_fport(ha, tq, tq->loop_id,
2347             LFF_NO_PLOGI, &mr);
2348         if (rval != QL_SUCCESS) {
2349             EL(ha, "failed, login_fport=%xh, "
2350                 "d_id=%xh\n", rval, tq->d_id.b24);
2351         }
2352     } else {
2353         rval = ql_login_lport(ha, tq, tq->loop_id,
2354             LLF_NONE);

```

```

2355         if (rval != QL_SUCCESS) {
2356             EL(ha, "failed, login_lport=%xh, "
2357                "d_id=%xh\n", rval, tq->d_id.b24);
2358         }
2359     } else {
2360         break;
2361     }
2362 }
2364     bzero((caddr_t)pkt, sizeof (ql_mbx_iocb_t));
2366 } while (retries--);
2368 if (sts.comp_status == CS_LOOP_DOWN_ABORT) {
2369     /* Cannot issue command now, maybe later */
2370     EL(ha, "failed, suspended\n");
2371     kmem_free(pkt, pkt_size);
2372     ql_free_dma_resource(ha, dma_mem);
2373     kmem_free(dma_mem, sizeof (dma_mem_t));
2374     cmd->Status = EXT_STATUS_SUSPENDED;
2375     cmd->ResponseLen = 0;
2376     return;
2377 }
2379 if (status != QL_SUCCESS) {
2380     /* Command error */
2381     EL(ha, "failed, I/O\n");
2382     kmem_free(pkt, pkt_size);
2383     ql_free_dma_resource(ha, dma_mem);
2384     kmem_free(dma_mem, sizeof (dma_mem_t));
2385     cmd->Status = EXT_STATUS_ERR;
2386     cmd->DetailStatus = status;
2387     cmd->ResponseLen = 0;
2388     return;
2389 }
2391 /* Setup status. */
2392 if (CFG_IST(ha, CFG_CTRL_24258081)) {
2393     sts.scsi_status_l = pkt->sts24.scsi_status_l;
2394     sts.scsi_status_h = pkt->sts24.scsi_status_h;
2396     /* Setup residuals. */
2397     sts.residual_length = LE_32(pkt->sts24.residual_length);
2399     /* Setup state flags. */
2400     sts.state_flags_l = pkt->sts24.state_flags_l;
2401     sts.state_flags_h = pkt->sts24.state_flags_h;
2402     if (pld_size && sts.comp_status != CS_DATA_UNDERRUN) {
2403         sts.state_flags_h = (uint8_t)(sts.state_flags_h |
2404             SF_GOT_BUS | SF_GOT_TARGET | SF_SENT_CMD |
2405             SF_XFERRED_DATA | SF_GOT_STATUS);
2406     } else {
2407         sts.state_flags_h = (uint8_t)(sts.state_flags_h |
2408             SF_GOT_BUS | SF_GOT_TARGET | SF_SENT_CMD |
2409             SF_GOT_STATUS);
2410     }
2411     if (scsi_req.direction & CF_WR) {
2412         sts.state_flags_l = (uint8_t)(sts.state_flags_l |
2413             SF_DATA_OUT);
2414     } else if (scsi_req.direction & CF_RD) {
2415         sts.state_flags_l = (uint8_t)(sts.state_flags_l |
2416             SF_DATA_IN);
2417     }
2418     sts.state_flags_l = (uint8_t)(sts.state_flags_l | SF_SIMPLE_Q);
2420     /* Setup FCP response info. */

```

```

2421     sts.rsp_info_length = sts.scsi_status_h & FCP_RSP_LEN_VALID ?
2422         LE_32(pkt->sts24.fcp_rsp_data_length) : 0;
2423     sts.rsp_info = &pkt->sts24.rsp_sense_data[0];
2424     for (cnt = 0; cnt < sts.rsp_info_length;
2425          cnt = (uint16_t)(cnt + 4)) {
2426         ql_chg_endian(sts.rsp_info + cnt, 4);
2427     }
2429     /* Setup sense data. */
2430     if (sts.scsi_status_h & FCP_SNS_LEN_VALID) {
2431         sts.req_sense_length =
2432             LE_32(pkt->sts24.fcp_sense_length);
2433         sts.state_flags_h = (uint8_t)(sts.state_flags_h |
2434             SF_ARQ_DONE);
2435     } else {
2436         sts.req_sense_length = 0;
2437     }
2438     sts.req_sense_data =
2439         &pkt->sts24.rsp_sense_data[sts.rsp_info_length];
2440     cnt2 = (uint16_t)((uintptr_t)pkt + sizeof (sts_24xx_entry_t)) -
2441         (uintptr_t)sts.req_sense_data;
2442     for (cnt = 0; cnt < cnt2; cnt = (uint16_t)(cnt + 4)) {
2443         ql_chg_endian(sts.req_sense_data + cnt, 4);
2444     }
2445 } else {
2446     sts.scsi_status_l = pkt->sts.scsi_status_l;
2447     sts.scsi_status_h = pkt->sts.scsi_status_h;
2449     /* Setup residuals. */
2450     sts.residual_length = LE_32(pkt->sts.residual_length);
2452     /* Setup state flags. */
2453     sts.state_flags_l = pkt->sts.state_flags_l;
2454     sts.state_flags_h = pkt->sts.state_flags_h;
2456     /* Setup FCP response info. */
2457     sts.rsp_info_length = sts.scsi_status_h & FCP_RSP_LEN_VALID ?
2458         LE_16(pkt->sts.rsp_info_length) : 0;
2459     sts.rsp_info = &pkt->sts.rsp_info[0];
2461     /* Setup sense data. */
2462     sts.req_sense_length = sts.scsi_status_h & FCP_SNS_LEN_VALID ?
2463         LE_16(pkt->sts.req_sense_length) : 0;
2464     sts.req_sense_data = &pkt->sts.req_sense_data[0];
2465 }
2467     QL_PRINT_9(CE_CONT, "(%d): response pkt\n", ha->instance);
2468     QL_DUMP_9(&pkt->sts, 8, sizeof (sts_entry_t));
2470     switch (sts.comp_status) {
2471     case CS_INCOMPLETE:
2472     case CS_ABORTED:
2473     case CS_DEVICE_UNAVAILABLE:
2474     case CS_PORT_UNAVAILABLE:
2475     case CS_PORT_LOGGED_OUT:
2476     case CS_PORT_CONFIG_CHG:
2477     case CS_PORT_BUSY:
2478     case CS_LOOP_DOWN_ABORT:
2479         cmd->Status = EXT_STATUS_BUSY;
2480         break;
2481     case CS_RESET:
2482     case CS_QUEUE_FULL:
2483         cmd->Status = EXT_STATUS_ERR;
2484         break;
2485     case CS_TIMEOUT:
2486         cmd->Status = EXT_STATUS_ERR;

```

```

2487         break;
2488     case CS_DATA_OVERRUN:
2489         cmd->Status = EXT_STATUS_DATA_OVERRUN;
2490         break;
2491     case CS_DATA_UNDERRUN:
2492         cmd->Status = EXT_STATUS_DATA_UNDERRUN;
2493         break;
2494     }
2495
2496     /*
2497     * If non data transfer commands fix tranfer counts.
2498     */
2499     if (scsi_req.cdbp[0] == SCMD_TEST_UNIT_READY ||
2500         scsi_req.cdbp[0] == SCMD_REZERO_UNIT ||
2501         scsi_req.cdbp[0] == SCMD_SEEK ||
2502         scsi_req.cdbp[0] == SCMD_SEEK_G1 ||
2503         scsi_req.cdbp[0] == SCMD_RESERVE ||
2504         scsi_req.cdbp[0] == SCMD_RELEASE ||
2505         scsi_req.cdbp[0] == SCMD_START_STOP ||
2506         scsi_req.cdbp[0] == SCMD_DOORLOCK ||
2507         scsi_req.cdbp[0] == SCMD_VERIFY ||
2508         scsi_req.cdbp[0] == SCMD_WRITE_FILE_MARK ||
2509         scsi_req.cdbp[0] == SCMD_VERIFY_G0 ||
2510         scsi_req.cdbp[0] == SCMD_SPACE ||
2511         scsi_req.cdbp[0] == SCMD_ERASE ||
2512         (scsi_req.cdbp[0] == SCMD_FORMAT &&
2513         (scsi_req.cdbp[1] & FPB_DATA) == 0)) {
2514         /*
2515         * Non data transfer command, clear sts_entry residual
2516         * length.
2517         */
2518         sts.residual_length = 0;
2519         cmd->ResponseLen = 0;
2520         if (sts.comp_status == CS_DATA_UNDERRUN) {
2521             sts.comp_status = CS_COMPLETE;
2522             cmd->Status = EXT_STATUS_OK;
2523         }
2524     } else {
2525         cmd->ResponseLen = pld_size;
2526     }
2527
2528     /* Correct ISP completion status */
2529     if (sts.comp_status == CS_COMPLETE && sts.scsi_status_l == 0 &&
2530         (sts.scsi_status_h & FCP_RSP_MASK) == 0) {
2531         QL_PRINT_9(CE_CONT, "(%d): Correct completion\n",
2532             ha->instance);
2533         scsi_req.resid = 0;
2534     } else if (sts.comp_status == CS_DATA_UNDERRUN) {
2535         QL_PRINT_9(CE_CONT, "(%d): Correct UNDERRUN\n",
2536             ha->instance);
2537         scsi_req.resid = sts.residual_length;
2538         if (sts.scsi_status_h & FCP_RESID_UNDER) {
2539             cmd->Status = (uint32_t)EXT_STATUS_OK;
2540
2541             cmd->ResponseLen = (uint32_t)
2542                 (pld_size - scsi_req.resid);
2543         } else {
2544             EL(ha, "failed, Transfer ERROR\n");
2545             cmd->Status = EXT_STATUS_ERR;
2546             cmd->ResponseLen = 0;
2547         }
2548     } else {
2549         QL_PRINT_9(CE_CONT, "(%d): error d_id=%xh, comp_status=%xh, "
2550             "scsi_status_h=%xh, scsi_status_l=%xh\n", ha->instance,
2551             tq->d_id.b24, sts.comp_status, sts.scsi_status_h,
2552             sts.scsi_status_l);

```

```

2554         scsi_req.resid = pld_size;
2555     /*
2556     * Handle residual count on SCSI check
2557     * condition.
2558     *
2559     * - If Residual Under / Over is set, use the
2560     *   Residual Transfer Length field in IOCB.
2561     * - If Residual Under / Over is not set, and
2562     *   Transferred Data bit is set in State Flags
2563     *   field of IOCB, report residual value of 0
2564     *   (you may want to do this for tape
2565     *   Write-type commands only). This takes care
2566     *   of logical end of tape problem and does
2567     *   not break Unit Attention.
2568     * - If Residual Under / Over is not set, and
2569     *   Transferred Data bit is not set in State
2570     *   Flags, report residual value equal to
2571     *   original data transfer length.
2572     */
2573     if (sts.scsi_status_l & STATUS_CHECK) {
2574         cmd->Status = EXT_STATUS_SCSI_STATUS;
2575         cmd->DetailStatus = sts.scsi_status_l;
2576         if (sts.scsi_status_h &
2577             (FCP_RESID_OVER | FCP_RESID_UNDER)) {
2578             scsi_req.resid = sts.residual_length;
2579         } else if (sts.state_flags_h &
2580             STATE_XFERRED_DATA) {
2581             scsi_req.resid = 0;
2582         }
2583     }
2584 }
2585
2586 if (sts.scsi_status_l & STATUS_CHECK &&
2587     sts.scsi_status_h & FCP_SNS_LEN_VALID &&
2588     sts.req_sense_length) {
2589     /*
2590     * Check condition with vaild sense data flag set and sense
2591     * length != 0
2592     */
2593     if (sts.req_sense_length > scsi_req.sense_length) {
2594         sense_sz = scsi_req.sense_length;
2595     } else {
2596         sense_sz = sts.req_sense_length;
2597     }
2598
2599     EL(ha, "failed, Check Condition Status, d_id=%xh\n",
2600         tq->d_id.b24);
2601     QL_DUMP_2(sts.req_sense_data, 8, sts.req_sense_length);
2602
2603     if (ddi_copyout(sts.req_sense_data, scsi_req.u_sense,
2604         (size_t)sense_sz, mode) != 0) {
2605         EL(ha, "failed, request sense ddi_copyout\n");
2606     }
2607
2608     cmd->Status = EXT_STATUS_SCSI_STATUS;
2609     cmd->DetailStatus = sts.scsi_status_l;
2610 }
2611
2612 /* Copy response payload from DMA buffer to application. */
2613 if (scsi_req.direction & (CF_RD | CF_DATA_IN) &&
2614     cmd->ResponseLen != 0) {
2615     QL_PRINT_9(CE_CONT, "(%d): Data Return resid=%lu, "
2616         "byte_count=%u, ResponseLen=%xh\n", ha->instance,
2617         scsi_req.resid, pld_size, cmd->ResponseLen);
2618     QL_DUMP_9(pld, 8, cmd->ResponseLen);

```

```

2620     /* Send response payload. */
2621     if (ql_send_buffer_data(pld,
2622         (caddr_t)(uintptr_t)cmd->ResponseAdr,
2623         cmd->ResponseLen, mode) != cmd->ResponseLen) {
2624         EL(ha, "failed, send_buffer_data\n");
2625         cmd->Status = EXT_STATUS_COPY_ERR;
2626         cmd->ResponseLen = 0;
2627     }
2628 }

2630 if (cmd->Status != EXT_STATUS_OK) {
2631     EL(ha, "failed, cmd->Status=%xh, comp_status=%xh, "
2632         "d_id=%xh\n", cmd->Status, sts.comp_status, tq->d_id.b24);
2633 } else {
2634     /*EMPTY*/
2635     QL_PRINT_9(CE_CONT, "(%d): done, ResponseLen=%d\n",
2636         ha->instance, cmd->ResponseLen);
2637 }

2639 kmem_free(pkt, pkt_size);
2640 ql_free_dma_resource(ha, dma_mem);
2641 kmem_free(dma_mem, sizeof(dma_mem_t));
2642 }

```

unchanged portion omitted

```

3288 /*
3289 * ql_diagnostic_loopback
3290 * Performs EXT_CC_LOOPBACK Command
3291 *
3292 * Input:
3293 *   ha: adapter state pointer.
3294 *   cmd: Local EXT_IOCTL cmd struct pointer.
3295 *   mode: flags.
3296 *
3297 * Returns:
3298 *   None, request status indicated in cmd->Status.
3299 *
3300 * Context:
3301 *   Kernel context.
3302 */
3303 static void
3304 ql_diagnostic_loopback(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
3305 {
3306     EXT_LOOPBACK_REQ    plbreq;
3307     EXT_LOOPBACK_RSP    plbrsp;
3308     ql_mbx_data_t       mr;
3309     uint32_t            rval;
3310     caddr_t             bp;
3311     uint16_t            opt;

3313     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

3315     /* Get loop back request. */
3316     if (ddi_copyin((void *) (uintptr_t) cmd->RequestAdr,
3317         (void *)&plbreq, sizeof(EXT_LOOPBACK_REQ), mode) != 0) {
3318         EL(ha, "failed, ddi_copyin\n");
3319         cmd->Status = EXT_STATUS_COPY_ERR;
3320         cmd->ResponseLen = 0;
3321         return;
3322     }

3324     opt = (uint16_t)(plbreq.Options & MBC_LOOPBACK_POINT_MASK);

3326     /* Check transfer length fits in buffer. */
3327     if (plbreq.BufferLength < plbreq.TransferCount &&

```

```

3328     plbreq.TransferCount < MAILBOX_BUFFER_SIZE) {
3329         EL(ha, "failed, BufferLength=%d, xfercnt=%d, "
3330             "mailbox_buffer_size=%d\n", plbreq.BufferLength,
3331             plbreq.TransferCount, MAILBOX_BUFFER_SIZE);
3332         cmd->Status = EXT_STATUS_INVALID_PARAM;
3333         cmd->ResponseLen = 0;
3334         return;
3335     }

3337     /* Allocate command memory. */
3338     bp = kmem_zalloc(plbreq.TransferCount, KM_SLEEP);
3339     if (bp == NULL) {
3340         EL(ha, "failed, kmem_zalloc\n");
3341         cmd->Status = EXT_STATUS_NO_MEMORY;
3342         cmd->ResponseLen = 0;
3343         return;
3344     }

3346     /* Get loopback data. */
3347     if (ql_get_buffer_data((caddr_t)(uintptr_t)plbreq.BufferAddress,
3348         bp, plbreq.TransferCount, mode) != plbreq.TransferCount) {
3349         EL(ha, "failed, ddi_copyin-2\n");
3350         kmem_free(bp, plbreq.TransferCount);
3351         cmd->Status = EXT_STATUS_COPY_ERR;
3352         cmd->ResponseLen = 0;
3353         return;
3354     }

3356     if ((ha->task_daemon_flags & (QL_LOOP_TRANSITION | DRIVER_STALL)) ||
3357         ql_stall_driver(ha, 0) != QL_SUCCESS) {
3358         EL(ha, "failed, LOOP_NOT_READY\n");
3359         kmem_free(bp, plbreq.TransferCount);
3360         cmd->Status = EXT_STATUS_BUSY;
3361         cmd->ResponseLen = 0;
3362         return;
3363     }

3365     /* Shutdown IP. */
3366     if (ha->flags & IP_INITIALIZED) {
3367         (void) ql_shutdown_ip(ha);
3368     }

3370     /* determine topology so we can send the loopback or the echo */
3371     /* Echo is supported on 2300's only and above */

3373     if (CFG_IST(ha, CFG_CTRL_8081)) {
3374         if (!(ha->task_daemon_flags & LOOP_DOWN) && opt ==
3375             MBC_LOOPBACK_POINT_EXTERNAL) {
3376             if (plbreq.TransferCount > 252) {
3377                 EL(ha, "transfer count (%d) > 252\n",
3378                     plbreq.TransferCount);
3379                 kmem_free(bp, plbreq.TransferCount);
3380                 cmd->Status = EXT_STATUS_INVALID_PARAM;
3381                 cmd->ResponseLen = 0;
3382                 return;
3383             }
3384             plbrsp.CommandSent = INT_DEF_LB_ECHO_CMD;
3385             rval = ql_diag_echo(ha, 0, bp, plbreq.TransferCount,
3386                 MBC_ECHO_ELS, &mr);
3387         } else {
3388             if (CFG_IST(ha, CFG_CTRL_81XX)) {
3389                 (void) ql_set_loop_point(ha, opt);
3390             }
3391             plbrsp.CommandSent = INT_DEF_LB_LOOPBACK_CMD;
3392             rval = ql_diag_loopback(ha, 0, bp, plbreq.TransferCount,
3393                 opt, plbreq.IterationCount, &mr);

```



```

3388         if (CFG_IST(ha, CFG_CTRL_81XX)) {
3389             (void) ql_set_loop_point(ha, 0);
3390         }
3391     } else {
3392     }
3393     if (!(ha->task_daemon_flags & LOOP_DOWN) &&
3394         (ha->topology & QL_F_PORT) &&
3395         ha->device_id >= 0x2300) {
3396         QL_PRINT_9(CE_CONT, "(%d): F_PORT topology -- using "
3397             "echo\n", ha->instance);
3398         plbrsp.CommandSent = INT_DEF_LB_ECHO_CMD;
3399         rval = ql_diag_echo(ha, 0, bp, plbreq.TransferCount,
3400             (uint16_t)(CFG_IST(ha, CFG_CTRL_8081) ?
3401                 MBC_ECHO_ELS : MBC_ECHO_64BIT), &mr);
3402     } else {
3403         plbrsp.CommandSent = INT_DEF_LB_LOOPBACK_CMD;
3404         rval = ql_diag_loopback(ha, 0, bp, plbreq.TransferCount,
3405             opt, plbreq.IterationCount, &mr);
3406     }
3407 }
3409 ql_restart_driver(ha);
3411 /* Restart IP if it was shutdown. */
3412 if (ha->flags & IP_ENABLED && !(ha->flags & IP_INITIALIZED)) {
3413     (void) ql_initialize_ip(ha);
3414     ql_isp_rcvbuf(ha);
3415 }
3417 if (rval != QL_SUCCESS) {
3418     EL(ha, "failed, diagnostic_loopback_mbx=%x\n", rval);
3419     kmem_free(bp, plbreq.TransferCount);
3420     cmd->Status = EXT_STATUS_MAILBOX;
3421     cmd->DetailStatus = rval;
3422     cmd->ResponseLen = 0;
3423     return;
3424 }
3426 /* Return loopback data. */
3427 if (ql_send_buffer_data(bp, (caddr_t)(uintptr_t)plbreq.BufferAddress,
3428     plbreq.TransferCount, mode) != plbreq.TransferCount) {
3429     EL(ha, "failed, ddi_copyout\n");
3430     kmem_free(bp, plbreq.TransferCount);
3431     cmd->Status = EXT_STATUS_COPY_ERR;
3432     cmd->ResponseLen = 0;
3433     return;
3434 }
3435 kmem_free(bp, plbreq.TransferCount);
3437 /* Return loopback results. */
3438 plbrsp.BufferAddress = plbreq.BufferAddress;
3439 plbrsp.BufferLength = plbreq.TransferCount;
3440 plbrsp.CompletionStatus = mr.mb[0];
3442 if (plbrsp.CommandSent == INT_DEF_LB_ECHO_CMD) {
3443     plbrsp.CrcErrorCount = 0;
3444     plbrsp.DisparityErrorCount = 0;
3445     plbrsp.FrameLengthErrorCount = 0;
3446     plbrsp.IterationCountLastError = 0;
3447 } else {
3448     plbrsp.CrcErrorCount = mr.mb[1];
3449     plbrsp.DisparityErrorCount = mr.mb[2];
3450     plbrsp.FrameLengthErrorCount = mr.mb[3];
3451     plbrsp.IterationCountLastError = (mr.mb[19] >> 16) | mr.mb[18];
3452 }

```

```

3454     rval = ddi_copyout((void *)&plbrsp,
3455         (void *) (uintptr_t)cmd->ResponseAdr,
3456         sizeof(EXT_LOOPBACK_RSP), mode);
3457     if (rval != 0) {
3458         EL(ha, "failed, ddi_copyout-2\n");
3459         cmd->Status = EXT_STATUS_COPY_ERR;
3460         cmd->ResponseLen = 0;
3461         return;
3462     }
3463     cmd->ResponseLen = sizeof(EXT_LOOPBACK_RSP);
3465     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
3466 }
    unchanged_portion_omitted
3534 /*
3535  * ql_send_els_rnid
3536  * IOCTL for extended link service RNID command.
3537  *
3538  * Input:
3539  *   ha: adapter state pointer.
3540  *   cmd: User space CT arguments pointer.
3541  *   mode: flags.
3542  *
3543  * Returns:
3544  *   None, request status indicated in cmd->Status.
3545  *
3546  * Context:
3547  *   Kernel context.
3548  */
3549 static void
3550 ql_send_els_rnid(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
3551 {
3552     EXT_RNID_REQ tmp_rnid;
3553     port_id_t tmp_fcid;
3554     caddr_t tmp_buf, bptr;
3555     uint32_t copy_len;
3556     ql_tgt_t *tg;
3557     EXT_RNID_DATA rnid_data;
3558     uint32_t loop_ready_wait = 10 * 60 * 10;
3559     int rval = 0;
3560     uint32_t local_hba = 0;
3562     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
3564     if (DRIVER_SUSPENDED(ha)) {
3565         EL(ha, "failed, LOOP_NOT_READY\n");
3566         cmd->Status = EXT_STATUS_BUSY;
3567         cmd->ResponseLen = 0;
3568         return;
3569     }
3571     if (cmd->RequestLen != sizeof(EXT_RNID_REQ)) {
3572         /* parameter error */
3573         EL(ha, "failed, RequestLen < EXT_RNID_REQ, Len=%x\n",
3574             cmd->RequestLen);
3575         cmd->Status = EXT_STATUS_INVALID_PARAM;
3576         cmd->DetailStatus = EXT_DSTATUS_REQUEST_LEN;
3577         cmd->ResponseLen = 0;
3578         return;
3579     }
3581     if (ddi_copyin((void *) (uintptr_t)cmd->RequestAdr,
3582         &tmp_rnid, cmd->RequestLen, mode) != 0) {
3583         EL(ha, "failed, ddi_copyin\n");
3584         cmd->Status = EXT_STATUS_COPY_ERR;

```

```

3585         cmd->ResponseLen = 0;
3586         return;
3587     }

3589     /* Find loop ID of the device */
3590     if (tmp_rnid.Addr.Type == EXT_DEF_TYPE_WWNN) {
3591         bptr = CFG_IST(ha, CFG_CTRL_24258081) ?
3592             (caddr_t)&ha->init_ctrl_blk.cb24.node_name :
3593             (caddr_t)&ha->init_ctrl_blk.cb.node_name;
3594         if (bcmp((void *)bptr, (void *)tmp_rnid.Addr.FcAddr.WWNN,
3595             EXT_DEF_WWN_NAME_SIZE) == 0) {
3596             local_hba = 1;
3597         } else {
3598             tq = ql_find_port(ha,
3599                 (uint8_t *)tmp_rnid.Addr.FcAddr.WWNN, QLNT_NODE);
3600         }
3601     } else if (tmp_rnid.Addr.Type == EXT_DEF_TYPE_WWPN) {
3602         bptr = CFG_IST(ha, CFG_CTRL_24258081) ?
3603             (caddr_t)&ha->init_ctrl_blk.cb24.port_name :
3604             (caddr_t)&ha->init_ctrl_blk.cb.port_name;
3605         if (bcmp((void *)bptr, (void *)tmp_rnid.Addr.FcAddr.WWPN,
3606             EXT_DEF_WWN_NAME_SIZE) == 0) {
3607             local_hba = 1;
3608         } else {
3609             tq = ql_find_port(ha,
3610                 (uint8_t *)tmp_rnid.Addr.FcAddr.WWPN, QLNT_PORT);
3611         }
3612     } else if (tmp_rnid.Addr.Type == EXT_DEF_TYPE_PORTID) {
3613         /*
3614          * Copy caller's d_id to tmp space.
3615          */
3616         bcopy(&tmp_rnid.Addr.FcAddr.Id[1], tmp_fcid.r.d_id,
3617             EXT_DEF_PORTID_SIZE_ACTUAL);
3618         BIG_ENDIAN_24(&tmp_fcid.r.d_id[0]);

3620         if (bcmp((void *)&ha->d_id, (void *)tmp_fcid.r.d_id,
3621             EXT_DEF_PORTID_SIZE_ACTUAL) == 0) {
3622             local_hba = 1;
3623         } else {
3624             tq = ql_find_port(ha, (uint8_t *)tmp_fcid.r.d_id,
3625                 QLNT_PID);
3626         }
3627     }

3629     /* Allocate memory for command. */
3630     tmp_buf = kmem_zalloc(SEND_RNID_RSP_SIZE, KM_SLEEP);
3631     if (tmp_buf == NULL) {
3632         EL(ha, "failed, kmem_zalloc\n");
3633         cmd->Status = EXT_STATUS_NO_MEMORY;
3634         cmd->ResponseLen = 0;
3635         return;
3636     }

3637     if (local_hba) {
3638         rval = ql_get_rnid_params(ha, SEND_RNID_RSP_SIZE, tmp_buf);
3639         if (rval != QL_SUCCESS) {
3640             EL(ha, "failed, get_rnid_params_mbx=%xh\n", rval);
3641             kmem_free(tmp_buf, SEND_RNID_RSP_SIZE);
3642             cmd->Status = EXT_STATUS_ERR;
3643             cmd->ResponseLen = 0;
3644             return;
3645         }
3646     }

3647     /* Save gotten RNID data. */
3648     bcopy(tmp_buf, &rnid_data, sizeof (EXT_RNID_DATA));

```

```

3645         /* Now build the Send RNID response */
3646         tmp_buf[0] = (char)(EXT_DEF_RNID_DFORMAT_TOPO_DISC);
3647         tmp_buf[1] = (2 * EXT_DEF_WWN_NAME_SIZE);
3648         tmp_buf[2] = 0;
3649         tmp_buf[3] = sizeof (EXT_RNID_DATA);

3651         if (CFG_IST(ha, CFG_CTRL_24258081)) {
3652             bcopy(ha->init_ctrl_blk.cb24.port_name, &tmp_buf[4],
3653                 EXT_DEF_WWN_NAME_SIZE);
3654             bcopy(ha->init_ctrl_blk.cb24.node_name,
3655                 &tmp_buf[4 + EXT_DEF_WWN_NAME_SIZE],
3656                 EXT_DEF_WWN_NAME_SIZE);
3657         } else {
3658             bcopy(ha->init_ctrl_blk.cb.port_name, &tmp_buf[4],
3659                 EXT_DEF_WWN_NAME_SIZE);
3660             bcopy(ha->init_ctrl_blk.cb.node_name,
3661                 &tmp_buf[4 + EXT_DEF_WWN_NAME_SIZE],
3662                 EXT_DEF_WWN_NAME_SIZE);
3663         }

3665         bcopy((uint8_t *)&rnid_data,
3666             &tmp_buf[4 + 2 * EXT_DEF_WWN_NAME_SIZE],
3667             sizeof (EXT_RNID_DATA));
3668     } else {
3669         if (tq == NULL) {
3670             /* no matching device */
3671             EL(ha, "failed, device not found\n");
3672             kmem_free(tmp_buf, SEND_RNID_RSP_SIZE);
3673             cmd->Status = EXT_STATUS_DEV_NOT_FOUND;
3674             cmd->DetailStatus = EXT_DSTATUS_TARGET;
3675             cmd->ResponseLen = 0;
3676             return;
3677         }

3679         /* Send command */
3680         rval = ql_send_rnid_els(ha, tq->loop_id,
3681             (uint8_t)tmp_rnid.DataFormat, SEND_RNID_RSP_SIZE, tmp_buf);
3682         if (rval != QL_SUCCESS) {
3683             EL(ha, "failed, send_rnid_mbx=%xh, id=%xh\n",
3684                 rval, tq->loop_id);
3685             while (LOOP_NOT_READY(ha)) {
3686                 ql_delay(ha, 100000);
3687                 if (loop_ready_wait-- == 0) {
3688                     EL(ha, "failed, loop not ready\n");
3689                     cmd->Status = EXT_STATUS_ERR;
3690                     cmd->ResponseLen = 0;
3691                 }
3692             }
3693             rval = ql_send_rnid_els(ha, tq->loop_id,
3694                 (uint8_t)tmp_rnid.DataFormat, SEND_RNID_RSP_SIZE,
3695                 tmp_buf);
3696             if (rval != QL_SUCCESS) {
3697                 /* error */
3698                 EL(ha, "failed, send_rnid_mbx=%xh, id=%xh\n",
3699                     rval, tq->loop_id);
3700                 kmem_free(tmp_buf, SEND_RNID_RSP_SIZE);
3701                 cmd->Status = EXT_STATUS_ERR;
3702                 cmd->ResponseLen = 0;
3703                 return;
3704             }
3705         }
3706     }

3708     /* Copy the response */
3709     copy_len = (cmd->ResponseLen > SEND_RNID_RSP_SIZE) ?
3710         SEND_RNID_RSP_SIZE : cmd->ResponseLen;

```

```

3712     if (ql_send_buffer_data(tmp_buf, (caddr_t)(uintptr_t)cmd->ResponseAdr,
3713         copy_len, mode) != copy_len) {
3714         cmd->Status = EXT_STATUS_COPY_ERR;
3715         EL(ha, "failed, ddi_copyout\n");
3716     } else {
3717         cmd->ResponseLen = copy_len;
3718         if (copy_len < SEND_RNID_RSP_SIZE) {
3719             cmd->Status = EXT_STATUS_DATA_OVERRUN;
3720             EL(ha, "failed, EXT_STATUS_DATA_OVERRUN\n");
3721         } else if (cmd->ResponseLen > SEND_RNID_RSP_SIZE) {
3722             cmd->Status = EXT_STATUS_DATA_UNDERRUN;
3723             EL(ha, "failed, EXT_STATUS_DATA_UNDERRUN\n");
3724         } else {
3725             cmd->Status = EXT_STATUS_OK;
3726             QL_PRINT_9(CE_CONT, "(%d): done\n",
3727                 ha->instance);
3728         }
3729     }
3730 }
3731
3732     kmem_free(tmp_buf, SEND_RNID_RSP_SIZE);
3733 }

```

unchanged portion omitted

```

3961 /*
3962 * ql_report_lun
3963 *   Get numbers of LUNS using report LUN command.
3964 *
3965 * Input:
3966 *   ha:   adapter state pointer.
3967 *   q:    target queue pointer.
3968 *
3969 * Returns:
3970 *   Number of LUNs.
3971 *
3972 * Context:
3973 *   Kernel context.
3974 */
3975 static int
3976 ql_report_lun(ql_adapter_state_t *ha, ql_tgt_t *tq)
3977 {
3978     int             rval;
3979     uint8_t         retries;
3980     ql_mbx_iocb_t   *pkt;
3981     ql_rpt_lun_lst_t *rpt;
3982     dma_mem_t       dma_mem;
3983     uint32_t         pkt_size, cnt;
3984     uint16_t         comp_status;
3985     uint8_t         scsi_status_h, scsi_status_l, *reqs;
3986
3987     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
3988
3989     if (DRIVER_SUSPENDED(ha)) {
3990         EL(ha, "failed, LOOP_NOT_READY\n");
3991         return (0);
3992     }
3993
3994     pkt_size = sizeof(ql_mbx_iocb_t) + sizeof(ql_rpt_lun_lst_t);
3995     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
3996     if (pkt == NULL) {
3997         EL(ha, "failed, kmem_zalloc\n");
3998         return (0);
3999     }
4000     rpt = (ql_rpt_lun_lst_t *)((caddr_t)pkt + sizeof(ql_mbx_iocb_t));

```

```

3998     /* Get DMA memory for the IOCB */
3999     if (ql_get_dma_mem(ha, &dma_mem, sizeof(ql_rpt_lun_lst_t),
4000         LITTLE_ENDIAN_DMA, QL_DMA_RING_ALIGN) != QL_SUCCESS) {
4001         cmn_err(CE_WARN, "%s(%d): DMA memory "
4002             "alloc failed", QL_NAME, ha->instance);
4003         kmem_free(pkt, pkt_size);
4004         return (0);
4005     }
4006
4007     for (retries = 0; retries < 4; retries++) {
4008         if (CFG_IST(ha, CFG_CTRL_24258081)) {
4009             pkt->cmd24.entry_type = IOCB_CMD_TYPE_7;
4010             pkt->cmd24.entry_count = 1;
4011
4012             /* Set N_port handle */
4013             pkt->cmd24.n_port_hdl = (uint16_t)LE_16(tq->loop_id);
4014
4015             /* Set target ID */
4016             pkt->cmd24.target_id[0] = tq->d_id.b.al_pa;
4017             pkt->cmd24.target_id[1] = tq->d_id.b.area;
4018             pkt->cmd24.target_id[2] = tq->d_id.b.domain;
4019
4020             /* Set Virtual Port ID */
4021             pkt->cmd24.vp_index = ha->vp_index;
4022
4023             /* Set ISP command timeout. */
4024             pkt->cmd24.timeout = LE_16(15);
4025
4026             /* Load SCSI CDB */
4027             pkt->cmd24.scsi_cdb[0] = SCMD_REPORT_LUNS;
4028             pkt->cmd24.scsi_cdb[6] =
4029                 MSB(MSW(sizeof(ql_rpt_lun_lst_t)));
4030             pkt->cmd24.scsi_cdb[7] =
4031                 LSB(MSW(sizeof(ql_rpt_lun_lst_t)));
4032             pkt->cmd24.scsi_cdb[8] =
4033                 MSB(LSW(sizeof(ql_rpt_lun_lst_t)));
4034             pkt->cmd24.scsi_cdb[9] =
4035                 LSB(LSW(sizeof(ql_rpt_lun_lst_t)));
4036             for (cnt = 0; cnt < MAX_CMDSZ; cnt += 4) {
4037                 ql_chg_endian((uint8_t *)&pkt->cmd24.scsi_cdb
4038                     + cnt, 4);
4039             }
4040
4041             /* Set tag queue control flags */
4042             pkt->cmd24.task = TA_STAG;
4043
4044             /* Set transfer direction. */
4045             pkt->cmd24.control_flags = CF_RD;
4046
4047             /* Set data segment count. */
4048             pkt->cmd24.dseg_count = LE_16(1);
4049
4050             /* Load total byte count. */
4051             /* Load data descriptor. */
4052             pkt->cmd24.dseg_0_address[0] = (uint32_t)
4053                 LE_32(LSD(dma_mem.cookie.dmac_laddress));
4054             pkt->cmd24.dseg_0_address[1] = (uint32_t)
4055                 LE_32(MSD(dma_mem.cookie.dmac_laddress));
4056             pkt->cmd24.total_byte_count =
4057                 LE_32(sizeof(ql_rpt_lun_lst_t));
4058             pkt->cmd24.dseg_0_length =
4059                 LE_32(sizeof(ql_rpt_lun_lst_t));
4060         } else if (CFG_IST(ha, CFG_ENABLE_64BIT_ADDRESSING)) {
4061             pkt->cmd3.entry_type = IOCB_CMD_TYPE_3;
4062             pkt->cmd3.entry_count = 1;
4063             if (CFG_IST(ha, CFG_EXT_FW_INTERFACE)) {

```

```

4064         pkt->cmd3.target_l = LSB(tq->loop_id);
4065         pkt->cmd3.target_h = MSB(tq->loop_id);
4066     } else {
4067         pkt->cmd3.target_h = LSB(tq->loop_id);
4068     }
4069     pkt->cmd3.control_flags_l = CF_DATA_IN | CF_STAG;
4070     pkt->cmd3.timeout = LE_16(15);
4071     pkt->cmd3.dseg_count = LE_16(1);
4072     pkt->cmd3.scsi_cdb[0] = SCMD_REPORT_LUNS;
4073     pkt->cmd3.scsi_cdb[6] =
4074         MSB(MSW(sizeof (ql_rpt_lun_lst_t)));
4075     pkt->cmd3.scsi_cdb[7] =
4076         LSB(MSW(sizeof (ql_rpt_lun_lst_t)));
4077     pkt->cmd3.scsi_cdb[8] =
4078         MSB(LSW(sizeof (ql_rpt_lun_lst_t)));
4079     pkt->cmd3.scsi_cdb[9] =
4080         LSB(LSW(sizeof (ql_rpt_lun_lst_t)));
4081     pkt->cmd3.byte_count =
4082         LE_32(sizeof (ql_rpt_lun_lst_t));
4083     pkt->cmd3.dseg_0_address[0] = (uint32_t)
4084         LE_32(LSD(dma_mem.cookie.dmac_laddress));
4085     pkt->cmd3.dseg_0_address[1] = (uint32_t)
4086         LE_32(MSD(dma_mem.cookie.dmac_laddress));
4087     pkt->cmd3.dseg_0_length =
4088         LE_32(sizeof (ql_rpt_lun_lst_t));
4089 } else {
4090     pkt->cmd.entry_type = IOCB_CMD_TYPE_2;
4091     pkt->cmd.entry_count = 1;
4092     if (CFG_IST(ha, CFG_EXT_FW_INTERFACE)) {
4093         pkt->cmd.target_l = LSB(tq->loop_id);
4094         pkt->cmd.target_h = MSB(tq->loop_id);
4095     } else {
4096         pkt->cmd.target_h = LSB(tq->loop_id);
4097     }
4098     pkt->cmd.control_flags_l = CF_DATA_IN | CF_STAG;
4099     pkt->cmd.timeout = LE_16(15);
4100     pkt->cmd.dseg_count = LE_16(1);
4101     pkt->cmd.scsi_cdb[0] = SCMD_REPORT_LUNS;
4102     pkt->cmd.scsi_cdb[6] =
4103         MSB(MSW(sizeof (ql_rpt_lun_lst_t)));
4104     pkt->cmd.scsi_cdb[7] =
4105         LSB(MSW(sizeof (ql_rpt_lun_lst_t)));
4106     pkt->cmd.scsi_cdb[8] =
4107         MSB(LSW(sizeof (ql_rpt_lun_lst_t)));
4108     pkt->cmd.scsi_cdb[9] =
4109         LSB(LSW(sizeof (ql_rpt_lun_lst_t)));
4110     pkt->cmd.byte_count =
4111         LE_32(sizeof (ql_rpt_lun_lst_t));
4112     pkt->cmd.dseg_0_address = (uint32_t)
4113         LE_32(LSD(dma_mem.cookie.dmac_laddress));
4114     pkt->cmd.dseg_0_length =
4115         LE_32(sizeof (ql_rpt_lun_lst_t));
4116 }

4118 rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt,
4119     sizeof (ql_mbx_iocb_t));

4121 /* Sync in coming DMA buffer. */
4122 (void) ddi_dma_sync(dma_mem.dma_handle, 0, dma_mem.size,
4123     DDI_DMA_SYNC_FORKERNEL);
4124 /* Copy in coming DMA data. */
4125 ddi_rep_get8(dma_mem.acc_handle, (uint8_t *)rpt,
4126     (uint8_t *)dma_mem.bp, dma_mem.size, DDI_DEV_AUTOINCR);

4128 if (CFG_IST(ha, CFG_CTRL_24258081)) {
4129     pkt->sts24.entry_status = (uint8_t)

```

```

4130         (pkt->sts24.entry_status & 0x3c);
4131         comp_status = (uint16_t)LE_16(pkt->sts24.comp_status);
4132         scsi_status_h = pkt->sts24.scsi_status_h;
4133         scsi_status_l = pkt->sts24.scsi_status_l;
4134         cnt = scsi_status_h & FCP_RSP_LEN_VALID ?
4135             LE_32(pkt->sts24.fcp_rsp_data_length) : 0;
4136         reqs = &pkt->sts24.rsp_sense_data[cnt];
4137     } else {
4138         pkt->sts.entry_status = (uint8_t)
4139             (pkt->sts.entry_status & 0x7e);
4140         comp_status = (uint16_t)LE_16(pkt->sts.comp_status);
4141         scsi_status_h = pkt->sts.scsi_status_h;
4142         scsi_status_l = pkt->sts.scsi_status_l;
4143         reqs = &pkt->sts.req_sense_data[0];
4144     }
4145     if (rval == QL_SUCCESS && pkt->sts.entry_status != 0) {
4146         EL(ha, "failed, entry_status=%xh, d_id=%xh\n",
4147             pkt->sts.entry_status, tq->d_id.b24);
4148         rval = QL_FUNCTION_PARAMETER_ERROR;
4149     }

4151     if (rval != QL_SUCCESS || comp_status != CS_COMPLETE ||
4152         scsi_status_l & STATUS_CHECK) {
4153         /* Device underrun, treat as OK. */
4154         if (rval == QL_SUCCESS &&
4155             comp_status == CS_DATA_UNDERRUN &&
4156             scsi_status_h & FCP_RESID_UNDER) {
4157             break;
4158         }

4160         EL(ha, "failed, issue_iocb=%xh, d_id=%xh, cs=%xh, "
4161             "ss_h=%xh, ss_l=%xh\n", rval, tq->d_id.b24,
4162             comp_status, scsi_status_h, scsi_status_l);

4164         if (rval == QL_SUCCESS) {
4165             if ((comp_status == CS_TIMEOUT) ||
4166                 (comp_status == CS_PORT_UNAVAILABLE) ||
4167                 (comp_status == CS_PORT_LOGGED_OUT)) {
4168                 rval = QL_FUNCTION_TIMEOUT;
4169                 break;
4170             }
4171             rval = QL_FUNCTION_FAILED;
4172         } else if (rval == QL_ABORTED) {
4173             break;
4174         }

4176         if (scsi_status_l & STATUS_CHECK) {
4177             EL(ha, "STATUS_CHECK Sense Data\n%2xh%3xh"
4178                 "%3xh%3xh%3xh%3xh%3xh%3xh%3xh%3xh%3xh"
4179                 "%3xh%3xh%3xh%3xh%3xh%3xh%3xh%3xh\n", reqs[0],
4180                 reqs[1], reqs[2], reqs[3], reqs[4],
4181                 reqs[5], reqs[6], reqs[7], reqs[8],
4182                 reqs[9], reqs[10], reqs[11], reqs[12],
4183                 reqs[13], reqs[14], reqs[15], reqs[16],
4184                 reqs[17]);
4185         }
4186     } else {
4187         break;
4188     }
4189     bzero((caddr_t)pkt, pkt_size);
4190 }

4192     if (rval != QL_SUCCESS) {
4193         EL(ha, "failed=%xh\n", rval);
4194         rval = 0;
4195     } else {

```

```

4196         QL_PRINT_9(CE_CONT, "(%d): LUN list\n", ha->instance);
4197         QL_DUMP_9(rpt, 8, rpt->hdr.len + 8);
4198         rval = (int)(BE_32(rpt->hdr.len) / 8);
4199     }

4201     kmem_free(pkt, pkt_size);
4202     ql_free_dma_resource(ha, &dma_mem);

4204     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

4206     return (rval);
4207 }

4209 /*
4210 * ql_inq_scan
4211 *   Get numbers of LUNS using inquiry command.
4212 *
4213 * Input:
4214 *   ha:          adapter state pointer.
4215 *   tq:          target queue pointer.
4216 *   count:      scan for the number of existing LUNs.
4217 *
4218 * Returns:
4219 *   Number of LUNs.
4220 *
4221 * Context:
4222 *   Kernel context.
4223 */
4224 static int
4225 ql_inq_scan(ql_adapter_state_t *ha, ql_tgt_t *tq, int count)
4226 {
4227     int             lun, cnt, rval;
4228     ql_mbx_iocb_t  *pkt;
4229     uint8_t        *inq;
4230     uint32_t       pkt_size;

4232     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

4234     pkt_size = sizeof (ql_mbx_iocb_t) + INQ_DATA_SIZE;
4235     pkt = kmem_zalloc(pkt_size, KM_SLEEP);
4236     if (pkt == NULL) {
4237         EL(ha, "failed, kmem_zalloc\n");
4238         return (0);
4239     }
4240     inq = (uint8_t *)((caddr_t)pkt + sizeof (ql_mbx_iocb_t));

4242     cnt = 0;
4243     for (lun = 0; lun < MAX_LUNS; lun++) {

4244         if (DRIVER_SUSPENDED(ha)) {
4245             rval = QL_LOOP_DOWN;
4246             cnt = 0;
4247             break;
4248         }

4249         rval = ql_inq(ha, tq, lun, pkt, INQ_DATA_SIZE);
4250         if (rval == QL_SUCCESS) {
4251             switch (*inq) {
4252                 case DTYPE_DIRECT:
4253                 case DTYPE_PROCESSOR: /* Appliance. */
4254                 case DTYPE_WORM:
4255                 case DTYPE_RODIRECT:
4256                 case DTYPE_SCANNER:
4257                 case DTYPE_OPTICAL:
4258                 case DTYPE_CHANGER:
4259                 case DTYPE_ESI:

```

```

4258         cnt++;
4259         break;
4260     case DTYPE_SEQUENTIAL:
4261         cnt++;
4262         tq->flags |= TQF_TAPE_DEVICE;
4263         break;
4264     default:
4265         QL_PRINT_9(CE_CONT, "(%d): failed, "
4266             "unsupported device id=%xh, lun=%d, "
4267             "type=%xh\n", ha->instance, tq->loop_id,
4268             lun, *inq);
4269         break;
4270     }

4272     if (*inq == DTYPE_ESI || cnt >= count) {
4273         break;
4274     }
4275     } else if (rval == QL_ABORTED || rval == QL_FUNCTION_TIMEOUT) {
4276         cnt = 0;
4277         break;
4278     }
4279 }

4281     kmem_free(pkt, pkt_size);

4283     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);

4285     return (cnt);
4286 }

unchanged_portion_omitted

4626 /*
4627 * ql_24xx_flash_desc
4628 *   Get flash descriptor table.
4629 *
4630 * Input:
4631 *   ha:          adapter state pointer.
4632 *
4633 * Returns:
4634 *   ql local function return status code.
4635 *
4636 * Context:
4637 *   Kernel context.
4638 */
4639 static int
4640 ql_24xx_flash_desc(ql_adapter_state_t *ha)
4641 {
4642     uint32_t        cnt;
4643     uint16_t        chksum, *bp, data;
4644     int             rval;
4645     flash_desc_t   *fdesc;
4646     ql_xioctl_t     *xp = ha->xioctl;

4648     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

4650     if (ha->flash_desc_addr == 0) {
4651         QL_PRINT_9(CE_CONT, "(%d): desc ptr=0\n", ha->instance);
4652         return (QL_FUNCTION_FAILED);
4653     }

4655     fdesc = kmem_zalloc(sizeof (flash_desc_t), KM_SLEEP);
4656     if ((fdesc = kmem_zalloc(sizeof (flash_desc_t), KM_SLEEP)) == NULL) {
4657         EL(ha, "kmem_zalloc=null\n");
4658         return (QL_MEMORY_ALLOC_FAILED);
4659     }
4660     rval = ql_dump_fcode(ha, (uint8_t *)fdesc, sizeof (flash_desc_t),

```

```

4657     ha->flash_desc_addr << 2);
4658     if (rval != QL_SUCCESS) {
4659         EL(ha, "read status=%xh\n", rval);
4660         kmem_free(fdesc, sizeof (flash_desc_t));
4661         return (rval);
4662     }
4664     chksum = 0;
4665     bp = (uint16_t *)fdesc;
4666     for (cnt = 0; cnt < (sizeof (flash_desc_t)) / 2; cnt++) {
4667         data = *bp++;
4668         LITTLE_ENDIAN_16(&data);
4669         chksum += data;
4670     }
4672     LITTLE_ENDIAN_32(&fdesc->flash_valid);
4673     LITTLE_ENDIAN_16(&fdesc->flash_version);
4674     LITTLE_ENDIAN_16(&fdesc->flash_len);
4675     LITTLE_ENDIAN_16(&fdesc->flash_checksum);
4676     LITTLE_ENDIAN_16(&fdesc->flash_manuf);
4677     LITTLE_ENDIAN_16(&fdesc->flash_id);
4678     LITTLE_ENDIAN_32(&fdesc->block_size);
4679     LITTLE_ENDIAN_32(&fdesc->alt_block_size);
4680     LITTLE_ENDIAN_32(&fdesc->flash_size);
4681     LITTLE_ENDIAN_32(&fdesc->write_enable_data);
4682     LITTLE_ENDIAN_32(&fdesc->read_timeout);
4684     /* flash size in desc table is in 1024 bytes */
4685     fdesc->flash_size = fdesc->flash_size * 0x400;
4687     if (chksum != 0 || fdesc->flash_valid != FLASH_DESC_VALID ||
4688         fdesc->flash_version != FLASH_DESC_VERSION) {
4689         EL(ha, "invalid descriptor table\n");
4690         kmem_free(fdesc, sizeof (flash_desc_t));
4691         return (QL_FUNCTION_FAILED);
4692     }
4694     bcopy(fdesc, &xp->fdesc, sizeof (flash_desc_t));
4695     kmem_free(fdesc, sizeof (flash_desc_t));
4697     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
4699     return (QL_SUCCESS);
4700 }

```

unchanged portion omitted

```

5342 /*
5343  * ql_set_rnid_parameters
5344  * Set RNID parameters.
5345  *
5346  * Input:
5347  * ha: adapter state pointer.
5348  * cmd: User space CT arguments pointer.
5349  * mode: flags.
5350  */
5351 static void
5352 ql_set_rnid_parameters(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
5353 {
5354     EXT_SET_RNID_REQ tmp_set;
5355     EXT_RNID_DATA *tmp_buf;
5356     int rval = 0;
5358     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
5360     if (DRIVER_SUSPENDED(ha)) {
5361         EL(ha, "failed, LOOP_NOT_READY\n");

```

```

5362     cmd->Status = EXT_STATUS_BUSY;
5363     cmd->ResponseLen = 0;
5364     return;
5365 }
5367     cmd->ResponseLen = 0; /* NO response to caller. */
5368     if (cmd->RequestLen != sizeof (EXT_SET_RNID_REQ)) {
5369         /* parameter error */
5370         EL(ha, "failed, RequestLen < EXT_SET_RNID_REQ, Len=%xh\n",
5371             cmd->RequestLen);
5372         cmd->Status = EXT_STATUS_INVALID_PARAM;
5373         cmd->DetailStatus = EXT_DSTATUS_REQUEST_LEN;
5374         cmd->ResponseLen = 0;
5375         return;
5376     }
5378     rval = ddi_copyin((void*)(uintptr_t)cmd->RequestAdr, &tmp_set,
5379         cmd->RequestLen, mode);
5380     if (rval != 0) {
5381         EL(ha, "failed, ddi_copyin\n");
5382         cmd->Status = EXT_STATUS_COPY_ERR;
5383         cmd->ResponseLen = 0;
5384         return;
5385     }
5387     /* Allocate memory for command. */
5388     tmp_buf = kmem_zalloc(sizeof (EXT_RNID_DATA), KM_SLEEP);
5389     if (tmp_buf == NULL) {
5390         EL(ha, "failed, kmem_zalloc\n");
5391         cmd->Status = EXT_STATUS_NO_MEMORY;
5392         cmd->ResponseLen = 0;
5393         return;
5394     }
5396     rval = ql_get_rnid_params(ha, sizeof (EXT_RNID_DATA),
5397         (caddr_t)tmp_buf);
5398     if (rval != QL_SUCCESS) {
5399         /* error */
5400         EL(ha, "failed, get_rnid_params_mbx=%xh\n", rval);
5401         kmem_free(tmp_buf, sizeof (EXT_RNID_DATA));
5402         cmd->Status = EXT_STATUS_ERR;
5403         cmd->ResponseLen = 0;
5404         return;
5405     }
5407     rval = ql_set_rnid_params(ha, sizeof (EXT_RNID_DATA),
5408         (caddr_t)tmp_buf);
5409     if (rval != QL_SUCCESS) {
5410         /* error */
5411         EL(ha, "failed, set_rnid_params_mbx=%xh\n", rval);
5412         cmd->Status = EXT_STATUS_ERR;
5413         cmd->ResponseLen = 0;
5414     }
5415     kmem_free(tmp_buf, sizeof (EXT_RNID_DATA));
5417     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
5418 }
5420 /*
5421  * ql_get_rnid_parameters

```

```

5422 *      Get RNID parameters.
5423 *
5424 * Input:
5425 *      ha:      adapter state pointer.
5426 *      cmd:     User space CT arguments pointer.
5427 *      mode:    flags.
5428 */
5429 static void
5430 ql_get_rnid_parameters(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
5431 {
5432     EXT_RNID_DATA    *tmp_buf;
5433     uint32_t         rval;
5434
5435     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
5436
5437     if (DRIVER_SUSPENDED(ha)) {
5438         EL(ha, "failed, LOOP_NOT_READY\n");
5439         cmd->Status = EXT_STATUS_BUSY;
5440         cmd->ResponseLen = 0;
5441         return;
5442     }
5443
5444     /* Allocate memory for command. */
5445     tmp_buf = kmem_zalloc(sizeof (EXT_RNID_DATA), KM_SLEEP);
5446     if (tmp_buf == NULL) {
5447         EL(ha, "failed, kmem_zalloc\n");
5448         cmd->Status = EXT_STATUS_NO_MEMORY;
5449         cmd->ResponseLen = 0;
5450         return;
5451     }
5452
5453     /* Send command */
5454     rval = ql_get_rnid_params(ha, sizeof (EXT_RNID_DATA),
5455                             (caddr_t)tmp_buf);
5456     if (rval != QL_SUCCESS) {
5457         /* error */
5458         EL(ha, "failed, get_rnid_params_mbx=%x\n", rval);
5459         kmem_free(tmp_buf, sizeof (EXT_RNID_DATA));
5460         cmd->Status = EXT_STATUS_ERR;
5461         cmd->ResponseLen = 0;
5462         return;
5463     }
5464
5465     /* Copy the response */
5466     if (ql_send_buffer_data((caddr_t)tmp_buf,
5467                           (caddr_t)(uintptr_t)cmd->ResponseAdr,
5468                           sizeof (EXT_RNID_DATA), mode) != sizeof (EXT_RNID_DATA)) {
5469         EL(ha, "failed, ddi_copyout\n");
5470         cmd->Status = EXT_STATUS_COPY_ERR;
5471         cmd->ResponseLen = 0;
5472     } else {
5473         QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
5474         cmd->ResponseLen = sizeof (EXT_RNID_DATA);
5475     }
5476
5477     kmem_free(tmp_buf, sizeof (EXT_RNID_DATA));
5478 }
5479
5480 unchanged portion omitted
5481
5482 /*
5483 * ql_get_statistics
5484 * Performs EXT_SC_GET_STATISTICS subcommand. of EXT_CC_GET_DATA.
5485 *
5486 * Input:
5487 *      ha:      adapter state pointer.
5488 *      cmd:     Local EXT_IOCTL cmd struct pointer.

```

```

5489 *      mode:    flags.
5490 *
5491 * Returns:
5492 *      None, request status indicated in cmd->Status.
5493 *
5494 * Context:
5495 *      Kernel context.
5496 */
5497 static void
5498 ql_get_statistics(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
5499 {
5500     EXT_HBA_PORT_STAT    ps = {0};
5501     ql_link_stats_t      *ls;
5502     int                   rval;
5503     ql_xioctl_t          *xp = ha->xioctl;
5504     int                   retry = 10;
5505
5506     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
5507
5508     while (ha->task_daemon_flags &
5509           (ABORT_ISP_ACTIVE | LOOP_RESYNC_ACTIVE | DRIVER_STALL)) {
5510         ql_delay(ha, 10000000); /* 10 second delay */
5511
5512         retry--;
5513
5514         if (retry == 0) { /* effectively 100 seconds */
5515             EL(ha, "failed, LOOP_NOT_READY\n");
5516             cmd->Status = EXT_STATUS_BUSY;
5517             cmd->ResponseLen = 0;
5518             return;
5519         }
5520     }
5521
5522     /* Allocate memory for command. */
5523     ls = kmem_zalloc(sizeof (ql_link_stats_t), KM_SLEEP);
5524     if (ls == NULL) {
5525         EL(ha, "failed, kmem_zalloc\n");
5526         cmd->Status = EXT_STATUS_NO_MEMORY;
5527         cmd->ResponseLen = 0;
5528         return;
5529     }
5530
5531     /*
5532      * I think these are supposed to be port statistics
5533      * the loop ID or port ID should be in cmd->Instance.
5534      */
5535     rval = ql_get_status_counts(ha, (uint16_t)
5536                               (ha->task_daemon_flags & LOOP_DOWN ? 0xFF : ha->loop_id),
5537                               sizeof (ql_link_stats_t), (caddr_t)ls, 0);
5538     if (rval != QL_SUCCESS) {
5539         EL(ha, "failed, get_link_status=%x, id=%x\n", rval,
5540             ha->loop_id);
5541         cmd->Status = EXT_STATUS_MAILBOX;
5542         cmd->DetailStatus = rval;
5543         cmd->ResponseLen = 0;
5544     } else {
5545         ps.ControllerErrorCount = xp->ControllerErrorCount;
5546         ps.DeviceErrorCount = xp->DeviceErrorCount;
5547         ps.IoCount = (uint32_t)(xp->IOInputRequests +
5548                               xp->IOOutputRequests + xp->IOControlRequests);
5549         ps.MBytesCount = (uint32_t)(xp->IOInputMByteCnt +
5550                                   xp->IOOutputMByteCnt);
5551         ps.LipResetCount = xp->TotalLipResets;
5552         ps.InterruptCount = xp->TotalInterrupts;
5553         ps.LinkFailureCount = LE_32(ls->link_fail_cnt);
5554         ps.LossOfSyncCount = LE_32(ls->sync_loss_cnt);

```

```

5602     ps.LossOfSignalsCount = LE_32(ls->signal_loss_cnt);
5603     ps.PrimitiveSeqProtocolErrorCount = LE_32(ls->prot_err_cnt);
5604     ps.InvalidTransmissionWordCount = LE_32(ls->inv_xmit_cnt);
5605     ps.InvalidCRCCount = LE_32(ls->inv_crc_cnt);

5607     rval = ddi_copyout((void *)&ps,
5608         (void *) (uintptr_t) cmd->ResponseAdr,
5609         sizeof (EXT_HBA_PORT_STAT), mode);
5610     if (rval != 0) {
5611         EL(ha, "failed, ddi_copyout\n");
5612         cmd->Status = EXT_STATUS_COPY_ERR;
5613         cmd->ResponseLen = 0;
5614     } else {
5615         cmd->ResponseLen = sizeof (EXT_HBA_PORT_STAT);
5616     }
5617 }

5619     kmem_free(ls, sizeof (ql_link_stats_t));

5621     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
5622 }

5624 /*
5625 * ql_get_statistics_fc
5626 *     Performs EXT_SC_GET_FC_STATISTICS subcommand. of EXT_CC_GET_DATA.
5627 *
5628 * Input:
5629 *     ha:     adapter state pointer.
5630 *     cmd:    Local EXT_IOCTL cmd struct pointer.
5631 *     mode:   flags.
5632 *
5633 * Returns:
5634 *     None, request status indicated in cmd->Status.
5635 *
5636 * Context:
5637 *     Kernel context.
5638 */
5639 static void
5640 ql_get_statistics_fc(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
5641 {
5642     EXT_HBA_PORT_STAT     ps = {0};
5643     ql_link_stats_t      *ls;
5644     int                   rval;
5645     uint16_t              qlnt;
5646     EXT_DEST_ADDR         pextdestaddr;
5647     uint8_t               *name;
5648     ql_tgt_t              *tq = NULL;
5649     int                   retry = 10;

5651     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);

5653     if (ddi_copyin((void *) (uintptr_t) cmd->RequestAdr,
5654         (void *) &pextdestaddr, sizeof (EXT_DEST_ADDR), mode) != 0) {
5655         EL(ha, "failed, ddi_copyin\n");
5656         cmd->Status = EXT_STATUS_COPY_ERR;
5657         cmd->ResponseLen = 0;
5658         return;
5659     }

5661     qlnt = QLNT_PORT;
5662     name = pextdestaddr.DestAddr.WWPN;

5664     QL_PRINT_9(CE_CONT, "%d): wwpn=%02x%02x%02x%02x%02x%02x%02x\n",
5665         ha->instance, name[0], name[1], name[2], name[3], name[4],
5666         name[5], name[6], name[7]);

```

```

5668     tq = ql_find_port(ha, name, qlnt);

5670     if (tq == NULL || !VALID_TARGET_ID(ha, tq->loop_id)) {
5671         EL(ha, "failed, fc_port not found\n");
5672         cmd->Status = EXT_STATUS_DEV_NOT_FOUND;
5673         cmd->ResponseLen = 0;
5674         return;
5675     }

5677     while (ha->task_daemon_flags &
5678         (ABORT_ISP_ACTIVE | LOOP_RESYNC_ACTIVE | DRIVER_STALL)) {
5679         ql_delay(ha, 1000000); /* 10 second delay */

5681         retry--;

5683         if (retry == 0) { /* effectively 100 seconds */
5684             EL(ha, "failed, LOOP_NOT_READY\n");
5685             cmd->Status = EXT_STATUS_BUSY;
5686             cmd->ResponseLen = 0;
5687             return;
5688         }

5689     }

5691     /* Allocate memory for command. */
5692     ls = kmem_zalloc(sizeof (ql_link_stats_t), KM_SLEEP);
5693     if (ls == NULL) {
5694         EL(ha, "failed, kmem_zalloc\n");
5695         cmd->Status = EXT_STATUS_NO_MEMORY;
5696         cmd->ResponseLen = 0;
5697         return;
5698     }

5699     rval = ql_get_link_status(ha, tq->loop_id, sizeof (ql_link_stats_t),
5700         (caddr_t) ls, 0);
5701     if (rval != QL_SUCCESS) {
5702         EL(ha, "failed, get_link_status=%xh, d_id=%xh\n", rval,
5703             tq->d_id.b24);
5704         cmd->Status = EXT_STATUS_MAILBOX;
5705         cmd->DetailStatus = rval;
5706         cmd->ResponseLen = 0;
5707     } else {
5708         ps.LinkFailureCount = LE_32(ls->link_fail_cnt);
5709         ps.LossOfSyncCount = LE_32(ls->sync_loss_cnt);
5710         ps.LossOfSignalsCount = LE_32(ls->signal_loss_cnt);
5711         ps.PrimitiveSeqProtocolErrorCount = LE_32(ls->prot_err_cnt);
5712         ps.InvalidTransmissionWordCount = LE_32(ls->inv_xmit_cnt);
5713         ps.InvalidCRCCount = LE_32(ls->inv_crc_cnt);

5714         rval = ddi_copyout((void *)&ps,
5715             (void *) (uintptr_t) cmd->ResponseAdr,
5716             sizeof (EXT_HBA_PORT_STAT), mode);

5718         if (rval != 0) {
5719             EL(ha, "failed, ddi_copyout\n");
5720             cmd->Status = EXT_STATUS_COPY_ERR;
5721             cmd->ResponseLen = 0;
5722         } else {
5723             cmd->ResponseLen = sizeof (EXT_HBA_PORT_STAT);
5724         }
5725     }

5726     kmem_free(ls, sizeof (ql_link_stats_t));

5727     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
5728 }

```

unchanged portion omitted


```

6878 /*
6879 * ql_flash_layout_table
6880 *   Obtains flash addresses from table
6881 *
6882 * Input:
6883 *   ha:          adapter state pointer.
6884 *   flt_paddr:  flash layout pointer address.
6885 *
6886 * Context:
6887 *   Kernel context.
6888 */
6889 static void
6890 ql_flash_layout_table(ql_adapter_state_t *ha, uint32_t flt_paddr)
6891 {
6892     ql_flt_ptr_t    *fptr;
6893     uint8_t         *bp;
6894     int             rval;
6895     uint32_t        len, faddr, cnt;
6896     uint16_t        chksum, w16;

6898     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);

6900     /* Process flash layout table header */
6901     len = sizeof (ql_flt_ptr_t);
6902     bp = kmem_zalloc(len, KM_SLEEP);
6903     if ((bp = kmem_zalloc(len, KM_SLEEP)) == NULL) {
6904         EL(ha, "kmem_zalloc=null\n");
6905         return;
6906     }

6907     /* Process pointer to flash layout table */
6908     if ((rval = ql_dump_fcode(ha, bp, len, flt_paddr)) != QL_SUCCESS) {
6909         EL(ha, "fptr dump_flash pos=%xh, status=%xh\n", flt_paddr,
6910             rval);
6911         kmem_free(bp, len);
6912         return;
6913     }
6914     fptr = (ql_flt_ptr_t *)bp;

6915     /* Verify pointer to flash layout table. */
6916     for (chksum = 0, cnt = 0; cnt < len; cnt += 2) {
6917         w16 = (uint16_t)CHAR_TO_SHORT(bp[cnt], bp[cnt + 1]);
6918         chksum += w16;
6919     }
6920     if (chksum != 0 || fptr->sig[0] != 'Q' || fptr->sig[1] != 'F' ||
6921         fptr->sig[2] != 'L' || fptr->sig[3] != 'T') {
6922         EL(ha, "ptr chksum=%xh, sig=%c%c%c%c\n", chksum, fptr->sig[0],
6923             fptr->sig[1], fptr->sig[2], fptr->sig[3]);
6924         kmem_free(bp, len);
6925         return;
6926     }
6927     faddr = CHAR_TO_LONG(fptr->addr[0], fptr->addr[1], fptr->addr[2],
6928         fptr->addr[3]);

6929     kmem_free(bp, len);

6930     ql_process_flt(ha, faddr);

6931     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
6932 }

6933 /*
6934 * ql_process_flt
6935 *   Obtains flash addresses from flash layout table
6936 */

```

```

6937 * Input:
6938 *   ha:          adapter state pointer.
6939 *   faddr:      flash layout table byte address.
6940 *
6941 * Context:
6942 *   Kernel context.
6943 */
6944 static void
6945 ql_process_flt(ql_adapter_state_t *ha, uint32_t faddr)
6946 {
6947     ql_flt_hdr_t    *fhdr;
6948     ql_flt_region_t *frgn;
6949     uint8_t         *bp, *eaddr, nv_rg, vpd_rg;
6950     int             rval;
6951     uint32_t        len, cnt, fe_addr;
6952     uint16_t        chksum, w16;

6953     QL_PRINT_9(CE_CONT, "%d): started faddr=%xh\n", ha->instance, faddr);

6954     /* Process flash layout table header */
6955     bp = kmem_zalloc(FLASH_LAYOUT_TABLE_SIZE, KM_SLEEP);
6956     if ((bp = kmem_zalloc(FLASH_LAYOUT_TABLE_SIZE, KM_SLEEP)) == NULL) {
6957         EL(ha, "kmem_zalloc=null\n");
6958         return;
6959     }
6960     fhdr = (ql_flt_hdr_t *)bp;

6961     /* Process flash layout table. */
6962     if ((rval = ql_dump_fcode(ha, bp, FLASH_LAYOUT_TABLE_SIZE, faddr)) !=
6963         QL_SUCCESS) {
6964         EL(ha, "fhdr dump_flash pos=%xh, status=%xh\n", faddr, rval);
6965         kmem_free(bp, FLASH_LAYOUT_TABLE_SIZE);
6966         return;
6967     }

6968     /* Verify flash layout table. */
6969     len = (uint32_t)(CHAR_TO_SHORT(fhdr->len[0], fhdr->len[1]) +
6970         sizeof (ql_flt_hdr_t) + sizeof (ql_flt_region_t));
6971     if (len > FLASH_LAYOUT_TABLE_SIZE) {
6972         chksum = 0xffff;
6973     } else {
6974         for (chksum = 0, cnt = 0; cnt < len; cnt += 2) {
6975             w16 = (uint16_t)CHAR_TO_SHORT(bp[cnt], bp[cnt + 1]);
6976             chksum += w16;
6977         }
6978     }
6979     w16 = CHAR_TO_SHORT(fhdr->version[0], fhdr->version[1]);
6980     if (chksum != 0 || w16 != 1) {
6981         EL(ha, "table chksum=%xh, version=%d\n", chksum, w16);
6982         kmem_free(bp, FLASH_LAYOUT_TABLE_SIZE);
6983         return;
6984     }
6985     eaddr = bp + len;

6986     /* Process Function/Port Configuration Map. */
6987     nv_rg = vpd_rg = 0;
6988     if (CFG_IST(ha, CFG_CTRL_8021)) {
6989         uint16_t     i;
6990         uint8_t      *mbp = eaddr;
6991         ql_fp_cfg_map_t *cmp = (ql_fp_cfg_map_t *)mbp;

6992         len = (uint32_t)(CHAR_TO_SHORT(cmp->hdr.len[0],
6993             cmp->hdr.len[1]));
6994         if (len > FLASH_LAYOUT_TABLE_SIZE) {
6995             chksum = 0xffff;
6996         } else {

```

```

7001         for (chksum = 0, cnt = 0; cnt < len; cnt += 2) {
7002             w16 = (uint16_t)CHAR_TO_SHORT(mbp[cnt],
7003                 mbp[cnt + 1]);
7004             chksum += w16;
7005         }
7006     }
7007     w16 = CHAR_TO_SHORT(cmp->hdr.version[0], cmp->hdr.version[1]);
7008     if (chksum != 0 || w16 != 1 ||
7009         cmp->hdr.Signature[0] != 'F' ||
7010         cmp->hdr.Signature[1] != 'P' ||
7011         cmp->hdr.Signature[2] != 'C' ||
7012         cmp->hdr.Signature[3] != 'M') {
7013         EL(ha, "cfg_map chksum=%xh, version=%d, "
7014             "sig=%c%c%c%c\n", chksum, w16,
7015             cmp->hdr.Signature[0], cmp->hdr.Signature[1],
7016             cmp->hdr.Signature[2], cmp->hdr.Signature[3]);
7017     } else {
7018         cnt = (uint16_t)
7019             (CHAR_TO_SHORT(cmp->hdr.NumberEntries[0],
7020                 cmp->hdr.NumberEntries[1]));
7021         /* Locate entry for function. */
7022         for (i = 0; i < cnt; i++) {
7023             if (cmp->cfg[i].FunctionType == FT_FC &&
7024                 cmp->cfg[i].FunctionNumber[0] ==
7025                 ha->function_number &&
7026                 cmp->cfg[i].FunctionNumber[1] == 0) {
7027                 nv_rg = cmp->cfg[i].ConfigRegion;
7028                 vpd_rg = cmp->cfg[i].VpdRegion;
7029                 break;
7030             }
7031         }
7032     }
7033     if (nv_rg == 0 || vpd_rg == 0) {
7034         EL(ha, "cfg_map nv_rg=%d, vpd_rg=%d\n", nv_rg,
7035             vpd_rg);
7036         nv_rg = vpd_rg = 0;
7037     }
7038 }
7039
7040 /* Process flash layout table regions */
7041 for (frgn = (qlflt_region_t *) (bp + sizeof(qlflt_hdr_t));
7042      (uint8_t *) frgn < eaddr; frgn++) {
7043     faddr = CHAR_TO_LONG(frgn->beg_addr[0], frgn->beg_addr[1],
7044         frgn->beg_addr[2], frgn->beg_addr[3]);
7045     faddr >>= 2;
7046     fe_addr = CHAR_TO_LONG(frgn->end_addr[0], frgn->end_addr[1],
7047         frgn->end_addr[2], frgn->end_addr[3]);
7048     fe_addr >>= 2;
7049
7050     switch (frgn->region) {
7051     case FLASH_8021_BOOTLOADER_REGION:
7052         ha->bootloader_addr = faddr;
7053         ha->bootloader_size = (fe_addr - faddr) + 1;
7054         QL_PRINT_9(CE_CONT, "(%d): bootloader_addr=%xh, "
7055             "size=%xh\n", ha->instance, faddr,
7056             ha->bootloader_size);
7057         break;
7058     case FLASH_FW_REGION:
7059     case FLASH_8021_FW_REGION:
7060         ha->flash_fw_addr = faddr;
7061         ha->flash_fw_size = (fe_addr - faddr) + 1;
7062         QL_PRINT_9(CE_CONT, "(%d): flash_fw_addr=%xh, "
7063             "size=%xh\n", ha->instance, faddr,
7064             ha->flash_fw_size);
7065         break;

```

```

7067         case FLASH_GOLDEN_FW_REGION:
7068     case FLASH_8021_GOLDEN_FW_REGION:
7069         ha->flash_golden_fw_addr = faddr;
7070         QL_PRINT_9(CE_CONT, "(%d): flash_golden_fw_addr=%xh\n",
7071             ha->instance, faddr);
7072         break;
7073     case FLASH_8021_VPD_REGION:
7074         if (!vpd_rg || vpd_rg == FLASH_8021_VPD_REGION) {
7075             ha->flash_vpd_addr = faddr;
7076             QL_PRINT_9(CE_CONT, "(%d): 8021_flash_vpd_"
7077                 "addr=%xh\n", ha->instance, faddr);
7078         }
7079         break;
7080     case FLASH_VPD_0_REGION:
7081         if (vpd_rg) {
7082             if (vpd_rg == FLASH_VPD_0_REGION) {
7083                 ha->flash_vpd_addr = faddr;
7084                 QL_PRINT_9(CE_CONT, "(%d): vpd_rg "
7085                     "flash_vpd_addr=%xh\n",
7086                     ha->instance, faddr);
7087             }
7088             } else if (!(ha->flags & FUNCTION_1) &&
7089                 !(CFG_IST(ha, CFG_CTRL_8021))) {
7090                 ha->flash_vpd_addr = faddr;
7091                 QL_PRINT_9(CE_CONT, "(%d): flash_vpd_addr=%xh"
7092                     "\n", ha->instance, faddr);
7093             }
7094         }
7095         break;
7096     case FLASH_NVRAM_0_REGION:
7097         if (nv_rg) {
7098             if (nv_rg == FLASH_NVRAM_0_REGION) {
7099                 ADAPTER_STATE_LOCK(ha);
7100                 ha->flags &= ~FUNCTION_1;
7101                 ADAPTER_STATE_UNLOCK(ha);
7102                 ha->flash_nvram_addr = faddr;
7103                 QL_PRINT_9(CE_CONT, "(%d): nv_rg "
7104                     "flash_nvram_addr=%xh\n",
7105                     ha->instance, faddr);
7106             }
7107             } else if (!(ha->flags & FUNCTION_1)) {
7108                 ha->flash_nvram_addr = faddr;
7109                 QL_PRINT_9(CE_CONT, "(%d): flash_nvram_addr="
7110                     "%xh\n", ha->instance, faddr);
7111             }
7112         }
7113         break;
7114     case FLASH_VPD_1_REGION:
7115         if (vpd_rg) {
7116             if (vpd_rg == FLASH_VPD_1_REGION) {
7117                 ha->flash_vpd_addr = faddr;
7118                 QL_PRINT_9(CE_CONT, "(%d): vpd_rg "
7119                     "flash_vpd_addr=%xh\n",
7120                     ha->instance, faddr);
7121             }
7122             } else if (ha->flags & FUNCTION_1 &&
7123                 !(CFG_IST(ha, CFG_CTRL_8021))) {
7124                 ha->flash_vpd_addr = faddr;
7125                 QL_PRINT_9(CE_CONT, "(%d): flash_vpd_addr=%xh"
7126                     "\n", ha->instance, faddr);
7127             }
7128         }
7129         break;
7130     case FLASH_NVRAM_1_REGION:
7131         if (nv_rg) {
7132             if (nv_rg == FLASH_NVRAM_1_REGION) {
7133                 ADAPTER_STATE_LOCK(ha);
7134                 ha->flags |= FUNCTION_1;
7135                 ADAPTER_STATE_UNLOCK(ha);

```

```

7133             ha->flash_nvram_addr = faddr;
7134             QL_PRINT_9(CE_CONT, "(%d): nv_rg "
7135                 "flash_nvram_addr=%xh\n",
7136                 ha->instance, faddr);
7137         }
7138     } else if (ha->flags & FUNCTION_1) {
7139         ha->flash_nvram_addr = faddr;
7140         QL_PRINT_9(CE_CONT, "(%d): flash_nvram_addr="
7141             "%xh\n", ha->instance, faddr);
7142     }
7143     break;
7144     case FLASH_DESC_TABLE_REGION:
7145         if (!(CFG_IST(ha, CFG_CTRL_8021))) {
7146             ha->flash_desc_addr = faddr;
7147             QL_PRINT_9(CE_CONT, "(%d): flash_desc_addr="
7148                 "%xh\n", ha->instance, faddr);
7149         }
7150     break;
7151     case FLASH_ERROR_LOG_0_REGION:
7152         if (!(ha->flags & FUNCTION_1)) {
7153             ha->flash_errlog_start = faddr;
7154             QL_PRINT_9(CE_CONT, "(%d): flash_errlog_addr="
7155                 "%xh\n", ha->instance, faddr);
7156         }
7157     break;
7158     case FLASH_ERROR_LOG_1_REGION:
7159         if (ha->flags & FUNCTION_1) {
7160             ha->flash_errlog_start = faddr;
7161             QL_PRINT_9(CE_CONT, "(%d): flash_errlog_addr="
7162                 "%xh\n", ha->instance, faddr);
7163         }
7164     break;
7165     default:
7166         break;
7167 }
7168 }
7169 kmem_free(bp, FLASH_LAYOUT_TABLE_SIZE);

7171     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
7172 }

```

unchanged portion omitted

```

7567 /*
7568 * ql_get_fwexttrace
7569 *     Dumps f/w extended trace buffer
7570 *
7571 * Input:
7572 *     ha:     adapter state pointer.
7573 *     bp:     buffer address.
7574 *     mode:   flags
7575 *
7576 * Returns:
7577 *
7578 * Context:
7579 *     Kernel context.
7580 */
7581 /* ARGSUSED */
7582 static void
7583 ql_get_fwexttrace(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
7584 {
7585     int rval;
7586     caddr_t payload;

7588     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

7590     if (CFG_IST(ha, CFG_CTRL_24258081) == 0) {

```

```

7591         EL(ha, "invalid request for this HBA\n");
7592         cmd->Status = EXT_STATUS_INVALID_REQUEST;
7593         cmd->ResponseLen = 0;
7594         return;
7595     }

7597     if ((CFG_IST(ha, CFG_ENABLE_FWEXTTRACE) == 0) ||
7598         (ha->fwexttracebuf.bp == NULL)) {
7599         EL(ha, "f/w extended trace is not enabled\n");
7600         cmd->Status = EXT_STATUS_INVALID_REQUEST;
7601         cmd->ResponseLen = 0;
7602         return;
7603     }

7605     if (cmd->ResponseLen < FWEXTSIZE) {
7606         cmd->Status = EXT_STATUS_BUFFER_TOO_SMALL;
7607         cmd->DetailStatus = FWEXTSIZE;
7608         EL(ha, "failed, ResponseLen (%xh) < %xh (FWEXTSIZE)\n",
7609             cmd->ResponseLen, FWEXTSIZE);
7610         cmd->ResponseLen = 0;
7611         return;
7612     }

7614     /* Time Stamp */
7615     rval = ql_fw_etrace(ha, &ha->fwexttracebuf, FTO_INSERT_TIME_STAMP);
7616     if (rval != QL_SUCCESS) {
7617         EL(ha, "f/w extended trace insert "
7618             "time stamp failed: %xh\n", rval);
7619         cmd->Status = EXT_STATUS_ERR;
7620         cmd->ResponseLen = 0;
7621         return;
7622     }

7624     /* Disable Tracing */
7625     rval = ql_fw_etrace(ha, &ha->fwexttracebuf, FTO_EXT_TRACE_DISABLE);
7626     if (rval != QL_SUCCESS) {
7627         EL(ha, "f/w extended trace disable failed: %xh\n", rval);
7628         cmd->Status = EXT_STATUS_ERR;
7629         cmd->ResponseLen = 0;
7630         return;
7631     }

7633     /* Allocate payload buffer */
7634     payload = kmem_zalloc(FWEXTSIZE, KM_SLEEP);
7635     if (payload == NULL) {
7636         EL(ha, "failed, kmem_zalloc\n");
7637         cmd->Status = EXT_STATUS_NO_MEMORY;
7638         cmd->ResponseLen = 0;
7639         return;
7640     }

7643     /* Sync DMA buffer. */
7644     (void) ddi_dma_sync(ha->fwexttracebuf.dma_handle, 0,
7645         FWEXTSIZE, DDI_DMA_SYNC_FORKERNEL);

7647     /* Copy trace buffer data. */
7648     ddi_rep_get8(ha->fwexttracebuf.acc_handle, (uint8_t *)payload,
7649         (uint8_t *)ha->fwexttracebuf.bp, FWEXTSIZE,
7650         DDI_DEV_AUTOINCR);

7652     /* Send payload to application. */
7653     if (ql_send_buffer_data(payload, (caddr_t)(uintptr_t)cmd->ResponseAdr,
7654         cmd->ResponseLen, mode) != cmd->ResponseLen) {
7655         EL(ha, "failed, send_buffer_data\n");
7656         cmd->Status = EXT_STATUS_COPY_ERR;
7657         cmd->ResponseLen = 0;

```

```

7651     } else {
7652         cmd->Status = EXT_STATUS_OK;
7653     }

7655     kmem_free(payload, FWEXTSIZE);

7657     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
7658 }

7660 /*
7661  * ql_get_fwfcetrace
7662  * Dumps f/w fibre channel event trace buffer
7663  *
7664  * Input:
7665  *   ha:   adapter state pointer.
7666  *   bp:   buffer address.
7667  *   mode: flags
7668  *
7669  * Returns:
7670  *
7671  * Context:
7672  *   Kernel context.
7673  */
7674 /* ARGSUSED */
7675 static void
7676 ql_get_fwfcetrace(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
7677 {
7678     int     rval;
7679     caddr_t payload;

7681     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);

7683     if (CFG_IST(ha, CFG_CTRL_24258081) == 0) {
7684         EL(ha, "invalid request for this HBA\n");
7685         cmd->Status = EXT_STATUS_INVALID_REQUEST;
7686         cmd->ResponseLen = 0;
7687         return;
7688     }

7690     if ((CFG_IST(ha, CFG_ENABLE_FWFCETRACE) == 0) ||
7691         (ha->fwfcetracebuf.bp == NULL)) {
7692         EL(ha, "f/w FCE trace is not enabled\n");
7693         cmd->Status = EXT_STATUS_INVALID_REQUEST;
7694         cmd->ResponseLen = 0;
7695         return;
7696     }

7698     if (cmd->ResponseLen < FWFCSIZE) {
7699         cmd->Status = EXT_STATUS_BUFFER_TOO_SMALL;
7700         cmd->DetailStatus = FWFCSIZE;
7701         EL(ha, "failed, ResponseLen (%xh) < %xh (FWFCSIZE)\n",
7702            cmd->ResponseLen, FWFCSIZE);
7703         cmd->ResponseLen = 0;
7704         return;
7705     }

7707     /* Disable Tracing */
7708     rval = ql_fw_etrace(ha, &ha->fwfcetracebuf, FTO_FCE_TRACE_DISABLE);
7709     if (rval != QL_SUCCESS) {
7710         EL(ha, "f/w FCE trace disable failed: %xh\n", rval);
7711         cmd->Status = EXT_STATUS_ERR;
7712         cmd->ResponseLen = 0;
7713         return;
7714     }

7716     /* Allocate payload buffer */

```

```

7717     payload = kmem_zalloc(FWEXTSIZE, KM_SLEEP);
7718     if (payload == NULL) {
7719         EL(ha, "failed, kmem_zalloc\n");
7720         cmd->Status = EXT_STATUS_NO_MEMORY;
7721         cmd->ResponseLen = 0;
7722         return;
7723     }

7719     /* Sync DMA buffer. */
7720     (void) ddi_dma_sync(ha->fwfcetracebuf.dma_handle, 0,
7721         FWFCSIZE, DDI_DMA_SYNC_FORKERNEL);

7723     /* Copy trace buffer data. */
7724     ddi_rep_get8(ha->fwfcetracebuf.acc_handle, (uint8_t *)payload,
7725         (uint8_t *)ha->fwfcetracebuf.bp, FWFCSIZE,
7726         DDI_DEV_AUTOINCR);

7728     /* Send payload to application. */
7729     if (ql_send_buffer_data(payload, (caddr_t)(uintptr_t)cmd->ResponseAdr,
7730         cmd->ResponseLen, mode) != cmd->ResponseLen) {
7731         EL(ha, "failed, send_buffer_data\n");
7732         cmd->Status = EXT_STATUS_COPY_ERR;
7733         cmd->ResponseLen = 0;
7734     } else {
7735         cmd->Status = EXT_STATUS_OK;
7736     }

7738     kmem_free(payload, FWFCSIZE);

7740     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
7741 }

    unchanged portion omitted

7812 /*
7813  * ql_pci_dump
7814  * Dumps PCI config data to application buffer.
7815  *
7816  * Input:
7817  *   ha = adapter state pointer.
7818  *   bp = user buffer address.
7819  *
7820  * Returns:
7821  *
7822  * Context:
7823  *   Kernel context.
7824  */
7825 int
7826 ql_pci_dump(ql_adapter_state_t *ha, uint32_t *bp, uint32_t pci_size, int mode)
7827 {
7828     uint32_t     pci_os;
7829     uint32_t     *ptr32, *org_ptr32;

7831     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);

7833     ptr32 = kmem_zalloc(pci_size, KM_SLEEP);
7834     if (ptr32 == NULL) {
7835         EL(ha, "failed kmem_zalloc\n");
7836         return (ENOMEM);
7837     }

7835     /* store the initial value of ptr32 */
7836     org_ptr32 = ptr32;
7837     for (pci_os = 0; pci_os < pci_size; pci_os += 4) {
7838         *ptr32 = (uint32_t)ql_pci_config_get32(ha, pci_os);
7839         LITTLE_ENDIAN_32(ptr32);
7840         ptr32++;

```

```

7841     }
7843     if (ddi_copyout((void *)org_ptr32, (void *)bp, pci_size, mode) !=
7844         0) {
7845         EL(ha, "failed ddi_copyout\n");
7846         kmem_free(org_ptr32, pci_size);
7847         return (EFAULT);
7848     }
7850     QL_DUMP_9(org_ptr32, 8, pci_size);
7852     kmem_free(org_ptr32, pci_size);
7854     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
7856     return (0);
7857 }

```

unchanged portion omitted

```

7940 /*
7941  * ql_menlo_get_fw_version
7942  *   Get Menlo firmware version.
7943  *
7944  * Input:
7945  *   ha:   adapter state pointer.
7946  *   bp:   buffer address.
7947  *   mode: flags
7948  *
7949  * Returns:
7950  *
7951  * Context:
7952  *   Kernel context.
7953  */
7954 static void
7955 ql_menlo_get_fw_version(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
7956 {
7957     int             rval;
7958     ql_mbx_iocb_t  *pkt;
7959     EXT_MENLO_GET_FW_VERSION ver = {0};
7961     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);
7963     if ((CFG_IST(ha, CFG_CTRL_MENLO)) == 0) {
7964         EL(ha, "failed, invalid request for HBA\n");
7965         cmd->Status = EXT_STATUS_INVALID_REQUEST;
7966         cmd->ResponseLen = 0;
7967         return;
7968     }
7970     if (cmd->ResponseLen < sizeof (EXT_MENLO_GET_FW_VERSION)) {
7971         cmd->Status = EXT_STATUS_BUFFER_TOO_SMALL;
7972         cmd->DetailStatus = sizeof (EXT_MENLO_GET_FW_VERSION);
7973         EL(ha, "ResponseLen=%d < %d\n", cmd->ResponseLen,
7974             sizeof (EXT_MENLO_GET_FW_VERSION));
7975         cmd->ResponseLen = 0;
7976         return;
7977     }
7979     /* Allocate packet. */
7980     pkt = kmem_zalloc(sizeof (ql_mbx_iocb_t), KM_SLEEP);
8076     if (pkt == NULL) {
8077         EL(ha, "failed, kmem_zalloc\n");
8078         cmd->Status = EXT_STATUS_NO_MEMORY;
8079         cmd->ResponseLen = 0;
8080         return;
8081     }

```

```

7982     pkt->mvfy.entry_type = VERIFY_MENLO_TYPE;
7983     pkt->mvfy.entry_count = 1;
7984     pkt->mvfy.options_status = LE_16(VMF_DO_NOT_UPDATE_FW);
7986     rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt, sizeof (ql_mbx_iocb_t));
7987     LITTLE_ENDIAN_16(&pkt->mvfy.options_status);
7988     LITTLE_ENDIAN_16(&pkt->mvfy.failure_code);
7989     ver.FwVersion = LE_32(pkt->mvfy.fw_version);
7991     if (rval != QL_SUCCESS || (pkt->mvfy.entry_status & 0x3c) != 0 ||
7992         pkt->mvfy.options_status != CS_COMPLETE) {
7993         /* Command error */
7994         EL(ha, "failed, status=%xh, es=%xh, cs=%xh, fc=%xh\n", rval,
7995             pkt->mvfy.entry_status & 0x3c, pkt->mvfy.options_status,
7996             pkt->mvfy.failure_code);
7997         cmd->Status = EXT_STATUS_ERR;
7998         cmd->DetailStatus = rval != QL_SUCCESS ? rval :
7999             QL_FUNCTION_FAILED;
8000         cmd->ResponseLen = 0;
8001     } else if (ddi_copyout((void *)&ver,
8002         (void *) (uintptr_t)cmd->ResponseAdr,
8003         sizeof (EXT_MENLO_GET_FW_VERSION), mode) != 0) {
8004         EL(ha, "failed, ddi_copyout\n");
8005         cmd->Status = EXT_STATUS_COPY_ERR;
8006         cmd->ResponseLen = 0;
8007     } else {
8008         cmd->ResponseLen = sizeof (EXT_MENLO_GET_FW_VERSION);
8009     }
8011     kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8013     QL_PRINT_9(CE_CONT, "%d): done\n", ha->instance);
8014 }
8016 /*
8017  * ql_menlo_update_fw
8018  *   Get Menlo update firmware.
8019  *
8020  * Input:
8021  *   ha:   adapter state pointer.
8022  *   bp:   buffer address.
8023  *   mode: flags
8024  *
8025  * Returns:
8026  *
8027  * Context:
8028  *   Kernel context.
8029  */
8030 static void
8031 ql_menlo_update_fw(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
8032 {
8033     ql_mbx_iocb_t  *pkt;
8034     dma_mem_t      *dma_mem;
8035     EXT_MENLO_UPDATE_FW fw;
8036     uint32_t       *ptr32;
8037     int             rval;
8039     QL_PRINT_9(CE_CONT, "%d): started\n", ha->instance);
8041     if ((CFG_IST(ha, CFG_CTRL_MENLO)) == 0) {
8042         EL(ha, "failed, invalid request for HBA\n");
8043         cmd->Status = EXT_STATUS_INVALID_REQUEST;
8044         cmd->ResponseLen = 0;
8045         return;
8046     }

```

```

8048 /*
8049  * TODO: only vp_index 0 can do this (?)
8050  */

8052 /* Verify the size of request structure. */
8053 if (cmd->RequestLen < sizeof (EXT_MENLO_UPDATE_FW)) {
8054     /* Return error */
8055     EL(ha, "RequestLen=%d < %d\n", cmd->RequestLen,
8056         sizeof (EXT_MENLO_UPDATE_FW));
8057     cmd->Status = EXT_STATUS_INVALID_PARAM;
8058     cmd->DetailStatus = EXT_DSTATUS_REQUEST_LEN;
8059     cmd->ResponseLen = 0;
8060     return;
8061 }

8063 /* Get update fw request. */
8064 if (ddi_copyin((caddr_t)(uintptr_t)cmd->RequestAdr, (caddr_t)&fw,
8065     sizeof (EXT_MENLO_UPDATE_FW), mode) != 0) {
8066     EL(ha, "failed, ddi_copyin\n");
8067     cmd->Status = EXT_STATUS_COPY_ERR;
8068     cmd->ResponseLen = 0;
8069     return;
8070 }

8072 /* Wait for I/O to stop and daemon to stall. */
8073 if (ql_suspend_hba(ha, 0) != QL_SUCCESS) {
8074     EL(ha, "ql_stall_driver failed\n");
8075     ql_restart_hba(ha);
8076     cmd->Status = EXT_STATUS_BUSY;
8077     cmd->ResponseLen = 0;
8078     return;
8079 }

8081 /* Allocate packet. */
8082 dma_mem = (dma_mem_t *)kmem_zalloc(sizeof (dma_mem_t), KM_SLEEP);
8083 if (dma_mem == NULL) {
8084     EL(ha, "failed, kmem_zalloc\n");
8085     cmd->Status = EXT_STATUS_NO_MEMORY;
8086     cmd->ResponseLen = 0;
8087     return;
8088 }
8089 pkt = kmem_zalloc(sizeof (ql_mbx_iocb_t), KM_SLEEP);
8090 if (pkt == NULL) {
8091     EL(ha, "failed, kmem_zalloc\n");
8092     kmem_free(dma_mem, sizeof (dma_mem_t));
8093     ql_restart_hba(ha);
8094     cmd->Status = EXT_STATUS_NO_MEMORY;
8095     cmd->ResponseLen = 0;
8096     return;
8097 }

8091 /* Get DMA memory for the IOCB */
8092 if (ql_get_dma_mem(ha, dma_mem, fw.TotalByteCount, LITTLE_ENDIAN_DMA,
8093     QL_DMA_DATA_ALIGN) != QL_SUCCESS) {
8094     cmn_err(CE_WARN, "%s(%d): request queue DMA memory "
8095         "alloc failed", QL_NAME, ha->instance);
8096     kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8097     kmem_free(dma_mem, sizeof (dma_mem_t));
8098     ql_restart_hba(ha);
8099     cmd->Status = EXT_STATUS_MS_NO_RESPONSE;
8100     cmd->ResponseLen = 0;
8101     return;
8102 }

8104 /* Get firmware data. */

```

```

8105     if (ql_get_buffer_data((caddr_t)(uintptr_t)fw.pFwDataBytes, dma_mem->bp,
8106         fw.TotalByteCount, mode) != fw.TotalByteCount) {
8107         EL(ha, "failed, get_buffer_data\n");
8108         ql_free_dma_resource(ha, dma_mem);
8109         kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8110         kmem_free(dma_mem, sizeof (dma_mem_t));
8111         ql_restart_hba(ha);
8112         cmd->Status = EXT_STATUS_COPY_ERR;
8113         cmd->ResponseLen = 0;
8114         return;
8115     }

8117 /* Sync DMA buffer. */
8118 (void) ddi_dma_sync(dma_mem->dma_handle, 0, dma_mem->size,
8119     DDI_DMA_SYNC_FORDEV);

8121 pkt->mvfy.entry_type = VERIFY_MENLO_TYPE;
8122 pkt->mvfy.entry_count = 1;
8123 pkt->mvfy.options_status = (uint16_t)LE_16(fw.Flags);
8124 ptr32 = dma_mem->bp;
8125 pkt->mvfy.fw_version = LE_32(ptr32[2]);
8126 pkt->mvfy.fw_size = LE_32(fw.TotalByteCount);
8127 pkt->mvfy.fw_sequence_size = LE_32(fw.TotalByteCount);
8128 pkt->mvfy.dseg_count = LE_16(1);
8129 pkt->mvfy.dseg_0_address[0] = (uint32_t)
8130     LE_32(LSD(dma_mem->cookie.dmac_laddress));
8131 pkt->mvfy.dseg_0_address[1] = (uint32_t)
8132     LE_32(MSD(dma_mem->cookie.dmac_laddress));
8133 pkt->mvfy.dseg_0_length = LE_32(fw.TotalByteCount);

8135 rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt, sizeof (ql_mbx_iocb_t));
8136 LITTLE_ENDIAN_16(&pkt->mvfy.options_status);
8137 LITTLE_ENDIAN_16(&pkt->mvfy.failure_code);

8139 if (rval != QL_SUCCESS || (pkt->mvfy.entry_status & 0x3c) != 0 ||
8140     pkt->mvfy.options_status != CS_COMPLETE) {
8141     /* Command error */
8142     EL(ha, "failed, status=%xh, es=%xh, cs=%xh, fc=%xh\n", rval,
8143         pkt->mvfy.entry_status & 0x3c, pkt->mvfy.options_status,
8144         pkt->mvfy.failure_code);
8145     cmd->Status = EXT_STATUS_ERR;
8146     cmd->DetailStatus = rval != QL_SUCCESS ? rval :
8147         QL_FUNCTION_FAILED;
8148     cmd->ResponseLen = 0;
8149 }

8151 ql_free_dma_resource(ha, dma_mem);
8152 kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8153 kmem_free(dma_mem, sizeof (dma_mem_t));
8154 ql_restart_hba(ha);

8156     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
8157 }

8159 /*
8160  * ql_menlo_manage_info
8161  * Get Menlo manage info.
8162  *
8163  * Input:
8164  * ha: adapter state pointer.
8165  * bp: buffer address.
8166  * mode: flags
8167  *
8168  * Returns:
8169  *
8170  * Context:

```

```

8171 *      Kernel context.
8172 */
8173 static void
8174 ql_menlo_manage_info(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
8175 {
8176     ql_mbx_iocb_t      *pkt;
8177     dma_mem_t          *dma_mem = NULL;
8178     EXT_MENLO_MANAGE_INFO info;
8179     int                 rval;
8181     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
8184     /* The call is only supported for Schultz right now */
8185     if (CFG_IST(ha, CFG_CTRL_8081)) {
8186         ql_get_xgmac_statistics(ha, cmd, mode);
8187         QL_PRINT_9(CE_CONT, "(%d): CFG_CTRL_81XX done\n",
8188                 ha->instance);
8189         return;
8190     }
8192     if (!CFG_IST(ha, CFG_CTRL_8081) || !CFG_IST(ha, CFG_CTRL_MENLO)) {
8193         EL(ha, "failed, invalid request for HBA\n");
8194         cmd->Status = EXT_STATUS_INVALID_REQUEST;
8195         cmd->ResponseLen = 0;
8196         return;
8197     }
8199     /* Verify the size of request structure. */
8200     if (cmd->RequestLen < sizeof (EXT_MENLO_MANAGE_INFO)) {
8201         /* Return error */
8202         EL(ha, "RequestLen=%d < %d\n", cmd->RequestLen,
8203             sizeof (EXT_MENLO_MANAGE_INFO));
8204         cmd->Status = EXT_STATUS_INVALID_PARAM;
8205         cmd->DetailStatus = EXT_DSTATUS_REQUEST_LEN;
8206         cmd->ResponseLen = 0;
8207         return;
8208     }
8210     /* Get manage info request. */
8211     if (ddi_copyin((caddr_t)(uintptr_t)cmd->RequestAdr,
8212                 (caddr_t)&info, sizeof (EXT_MENLO_MANAGE_INFO), mode) != 0) {
8213         EL(ha, "failed, ddi_copyin\n");
8214         cmd->Status = EXT_STATUS_COPY_ERR;
8215         cmd->ResponseLen = 0;
8216         return;
8217     }
8219     /* Allocate packet. */
8220     pkt = kmem_zalloc(sizeof (ql_mbx_iocb_t), KM_SLEEP);
8221     if (pkt == NULL) {
8222         EL(ha, "failed, kmem_zalloc\n");
8223         ql_restart_driver(ha);
8224         cmd->Status = EXT_STATUS_NO_MEMORY;
8225         cmd->ResponseLen = 0;
8226         return;
8227     }
8228     pkt->mdata.entry_type = MENLO_DATA_TYPE;
8229     pkt->mdata.entry_count = 1;
8230     pkt->mdata.options_status = (uint16_t)LE_16(info.Operation);
8232     /* Get DMA memory for the IOCB */
8233     if (info.Operation == MENLO_OP_READ_MEM ||
8234         info.Operation == MENLO_OP_WRITE_MEM) {
8235         pkt->mdata.total_byte_count = LE_32(info.TotalByteCount);

```

```

8230     pkt->mdata.parameter_1 =
8231         LE_32(info.Parameters.ap.MenloMemory.StartingAddr);
8232     dma_mem = (dma_mem_t *)kmem_zalloc(sizeof (dma_mem_t),
8233         KM_SLEEP);
8234     if (dma_mem == NULL) {
8235         EL(ha, "failed, kmem_zalloc\n");
8236         kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8237         cmd->Status = EXT_STATUS_NO_MEMORY;
8238         cmd->ResponseLen = 0;
8239         return;
8240     }
8241     if (ql_get_dma_mem(ha, dma_mem, info.TotalByteCount,
8242         LITTLE_ENDIAN_DMA, QL_DMA_DATA_ALIGN) != QL_SUCCESS) {
8243         cmn_err(CE_WARN, "%s(%d): request queue DMA memory "
8244             "alloc failed", QL_NAME, ha->instance);
8245         kmem_free(dma_mem, sizeof (dma_mem_t));
8246         kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8247         cmd->Status = EXT_STATUS_MS_NO_RESPONSE;
8248         cmd->ResponseLen = 0;
8249         return;
8250     }
8251     if (info.Operation == MENLO_OP_WRITE_MEM) {
8252         /* Get data. */
8253         if (ql_get_buffer_data(
8254             (caddr_t)(uintptr_t)info.pDataBytes,
8255             dma_mem->bp, info.TotalByteCount, mode) !=
8256             info.TotalByteCount) {
8257             EL(ha, "failed, get_buffer_data\n");
8258             ql_free_dma_resource(ha, dma_mem);
8259             kmem_free(dma_mem, sizeof (dma_mem_t));
8260             kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8261             cmd->Status = EXT_STATUS_COPY_ERR;
8262             cmd->ResponseLen = 0;
8263             return;
8264         }
8265         (void) ddi_dma_sync(dma_mem->dma_handle, 0,
8266             dma_mem->size, DDI_DMA_SYNC_FORDEV);
8267     }
8268     pkt->mdata.dseg_count = LE_16(1);
8269     pkt->mdata.dseg_0_address[0] = (uint32_t)
8270         LE_32(LSD(dma_mem->cookie.dmac_laddress));
8271     pkt->mdata.dseg_0_address[1] = (uint32_t)
8272         LE_32(MSD(dma_mem->cookie.dmac_laddress));
8273     pkt->mdata.dseg_0_length = LE_32(info.TotalByteCount);
8274     if (info.Operation & MENLO_OP_CHANGE_CONFIG) {
8275         pkt->mdata.parameter_1 =
8276             LE_32(info.Parameters.ap.MenloConfig.ConfigParamID);
8277         pkt->mdata.parameter_2 =
8278             LE_32(info.Parameters.ap.MenloConfig.ConfigParamData0);
8279         pkt->mdata.parameter_3 =
8280             LE_32(info.Parameters.ap.MenloConfig.ConfigParamData1);
8281     } else if (info.Operation & MENLO_OP_GET_INFO) {
8282         pkt->mdata.parameter_1 =
8283             LE_32(info.Parameters.ap.MenloInfo.InfoDataType);
8284         pkt->mdata.parameter_2 =
8285             LE_32(info.Parameters.ap.MenloInfo.InfoContext);
8286     }
8288     rval = ql_issue_mbx_iocb(ha, (caddr_t)pkt, sizeof (ql_mbx_iocb_t));
8289     LITTLE_ENDIAN_16(&pkt->mdata.options_status);
8290     LITTLE_ENDIAN_16(&pkt->mdata.failure_code);
8292     if (rval != QL_SUCCESS || (pkt->mdata.entry_status & 0x3c) != 0 ||
8293         pkt->mdata.options_status != CS_COMPLETE) {
8294         /* Command error */
8295         EL(ha, "failed, status=%xh, es=%xh, cs=%xh, fc=%xh\n", rval,

```

```

8296         pkt->mdata.entry_status & 0x3c, pkt->mdata.options_status,
8297         pkt->mdata.failure_code);
8298     cmd->Status = EXT_STATUS_ERR;
8299     cmd->DetailStatus = rval != QL_SUCCESS ? rval :
8300     QL_FUNCTION_FAILED;
8301     cmd->ResponseLen = 0;
8302 } else if (info.Operation == MENLO_OP_READ_MEM) {
8303     (void) ddi_dma_sync(dma_mem->dma_handle, 0, dma_mem->size,
8304     DDI_DMA_SYNC_FORKERNEL);
8305     if (ql_send_buffer_data((caddr_t)(uintptr_t)info.pDataBytes,
8306     dma_mem->bp, info.TotalByteCount, mode) !=
8307     info.TotalByteCount) {
8308         cmd->Status = EXT_STATUS_COPY_ERR;
8309         cmd->ResponseLen = 0;
8310     }
8311 }
8312
8313     ql_free_dma_resource(ha, dma_mem);
8314     kmem_free(dma_mem, sizeof (dma_mem_t));
8315     kmem_free(pkt, sizeof (ql_mbx_iocb_t));
8316
8317     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
8318 }
8319
8320 unchanged portion omitted
8321
8322 /*
8323  * ql_get_dcbx_parameters
8324  * Get DCBX parameters.
8325  *
8326  * Input:
8327  *   ha: adapter state pointer.
8328  *   cmd: User space CT arguments pointer.
8329  *   mode: flags.
8330  */
8331 static void
8332 ql_get_dcbx_parameters(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
8333 {
8334     uint8_t      *tmp_buf;
8335     int          rval;
8336
8337     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
8338
8339     if (!(CFG_IST(ha, CFG_CTRL_8081))) {
8340         EL(ha, "invalid request for HBA\n");
8341         cmd->Status = EXT_STATUS_INVALID_REQUEST;
8342         cmd->ResponseLen = 0;
8343         return;
8344     }
8345
8346     /* Allocate memory for command. */
8347     tmp_buf = kmem_zalloc(EXT_DEF_DCBX_PARAM_BUF_SIZE, KM_SLEEP);
8348     if (tmp_buf == NULL) {
8349         EL(ha, "failed, kmem_zalloc\n");
8350         cmd->Status = EXT_STATUS_NO_MEMORY;
8351         cmd->ResponseLen = 0;
8352         return;
8353     }
8354     /* Send command */
8355     rval = ql_get_dcbx_params(ha, EXT_DEF_DCBX_PARAM_BUF_SIZE,
8356     (caddr_t)tmp_buf);
8357     if (rval != QL_SUCCESS) {
8358         /* error */
8359         EL(ha, "failed, get_dcbx_params_mbx=%xh\n", rval);
8360         kmem_free(tmp_buf, EXT_DEF_DCBX_PARAM_BUF_SIZE);
8361         cmd->Status = EXT_STATUS_ERR;
8362         cmd->ResponseLen = 0;
8363     }
8364 }

```

```

8719         return;
8720     }
8721
8722     /* Copy the response */
8723     if (ql_send_buffer_data((caddr_t)tmp_buf,
8724     (caddr_t)(uintptr_t)cmd->ResponseAdr,
8725     EXT_DEF_DCBX_PARAM_BUF_SIZE, mode) != EXT_DEF_DCBX_PARAM_BUF_SIZE) {
8726         EL(ha, "failed, ddi_copyout\n");
8727         cmd->Status = EXT_STATUS_COPY_ERR;
8728         cmd->ResponseLen = 0;
8729     } else {
8730         cmd->ResponseLen = EXT_DEF_DCBX_PARAM_BUF_SIZE;
8731         QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
8732     }
8733     kmem_free(tmp_buf, EXT_DEF_DCBX_PARAM_BUF_SIZE);
8734 }
8735
8736 unchanged portion omitted
8737
8738 /*
8739  * ql_get_xgmac_statistics
8740  * Get XgMac information
8741  *
8742  * Input:
8743  *   ha: adapter state pointer.
8744  *   cmd: EXT_IOCTL cmd struct pointer.
8745  *   mode: flags.
8746  *
8747  * Returns:
8748  *   None, request status indicated in cmd->Status.
8749  *
8750  * Context:
8751  *   Kernel context.
8752  */
8753 static void
8754 ql_get_xgmac_statistics(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
8755 {
8756     int          rval;
8757     uint32_t    size;
8758     int8_t      *tmp_buf;
8759     EXT_MENLO_MANAGE_INFO info;
8760
8761     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);
8762
8763     /* Verify the size of request structure. */
8764     if (cmd->RequestLen < sizeof (EXT_MENLO_MANAGE_INFO)) {
8765         /* Return error */
8766         EL(ha, "RequestLen=%d < %d\n", cmd->RequestLen,
8767         sizeof (EXT_MENLO_MANAGE_INFO));
8768         cmd->Status = EXT_STATUS_INVALID_PARAM;
8769         cmd->DetailStatus = EXT_DSTATUS_REQUEST_LEN;
8770         cmd->ResponseLen = 0;
8771         return;
8772     }
8773
8774     /* Get manage info request. */
8775     if (ddi_copyin((caddr_t)(uintptr_t)cmd->RequestAdr,
8776     (caddr_t)&info, sizeof (EXT_MENLO_MANAGE_INFO), mode) != 0) {
8777         EL(ha, "failed, ddi_copyin\n");
8778         cmd->Status = EXT_STATUS_COPY_ERR;
8779         cmd->ResponseLen = 0;
8780         return;
8781     }
8782
8783     size = info.TotalByteCount;
8784     if (!size) {

```



```

8942         /* parameter error */
8943         cmd->Status = EXT_STATUS_INVALID_PARAM;
8944         cmd->DetailStatus = 0;
8945         EL(ha, "failed, size=%xh\n", size);
8946         cmd->ResponseLen = 0;
8947         return;
8948     }

8950     /* Allocate memory for command. */
8951     tmp_buf = kmem_zalloc(size, KM_SLEEP);
8952     if (tmp_buf == NULL) {
8953         EL(ha, "failed, kmem_zalloc\n");
8954         cmd->Status = EXT_STATUS_NO_MEMORY;
8955         cmd->ResponseLen = 0;
8956         return;
8957     }

8958     if (!(info.Operation & MENLO_OP_GET_INFO)) {
8959         EL(ha, "Invalid request for 81XX\n");
8960         kmem_free(tmp_buf, size);
8961         cmd->Status = EXT_STATUS_ERR;
8962         cmd->ResponseLen = 0;
8963         return;
8964     }

8965     rval = ql_get_xgmac_stats(ha, size, (caddr_t)tmp_buf);

8966     if (rval != QL_SUCCESS) {
8967         /* error */
8968         EL(ha, "failed, get_xgmac_stats=%xh\n", rval);
8969         kmem_free(tmp_buf, size);
8970         cmd->Status = EXT_STATUS_ERR;
8971         cmd->ResponseLen = 0;
8972         return;
8973     }

8974     if (ql_send_buffer_data(tmp_buf, (caddr_t)(uintptr_t)info.pDataBytes,
8975         size, mode) != size) {
8976         EL(ha, "failed, ddi_copyout\n");
8977         cmd->Status = EXT_STATUS_COPY_ERR;
8978         cmd->ResponseLen = 0;
8979     } else {
8980         cmd->ResponseLen = info.TotalByteCount;
8981         QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
8982     }
8983     kmem_free(tmp_buf, size);
8984     QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
8985 }

8986 /*
8987 * ql_get_fcf_list
8988 * Get FCF list.
8989 * Input:
8990 *   ha:      adapter state pointer.
8991 *   cmd:    User space CT arguments pointer.
8992 *   mode:   flags.
8993 */
8994 static void
8995 ql_get_fcf_list(ql_adapter_state_t *ha, EXT_IOCTL *cmd, int mode)
8996 {
8997     uint8_t      *tmp_buf;
8998     int          rval;
8999     EXT_FCF_LIST fcf_list = {0};
9000     ql_fcf_list_desc_t mb_fcf_list = {0};

```

```

9002     QL_PRINT_9(CE_CONT, "(%d): started\n", ha->instance);

9003     if (!(CFG_IST(ha, CFG_CTRL_81XX))) {
9004         EL(ha, "invalid request for HBA\n");
9005         cmd->Status = EXT_STATUS_INVALID_REQUEST;
9006         cmd->ResponseLen = 0;
9007         return;
9008     }

9009     /* Get manage info request. */
9010     if (ddi_copyin((caddr_t)(uintptr_t)cmd->RequestAdr,
9011         (caddr_t)&fcf_list, sizeof(EXT_FCF_LIST), mode) != 0) {
9012         EL(ha, "failed, ddi_copyin\n");
9013         cmd->Status = EXT_STATUS_COPY_ERR;
9014         cmd->ResponseLen = 0;
9015         return;
9016     }

9017     if (!(fcf_list.BufSize)) {
9018         /* Return error */
9019         EL(ha, "failed, fcf_list BufSize is=%xh\n",
9020             fcf_list.BufSize);
9021         cmd->Status = EXT_STATUS_INVALID_PARAM;
9022         cmd->ResponseLen = 0;
9023         return;
9024     }

9025     /* Allocate memory for command. */
9026     tmp_buf = kmem_zalloc(fcf_list.BufSize, KM_SLEEP);
9027     if (tmp_buf == NULL) {
9028         EL(ha, "failed, kmem_zalloc\n");
9029         cmd->Status = EXT_STATUS_NO_MEMORY;
9030         cmd->ResponseLen = 0;
9031         return;
9032     }

9033     /* build the descriptor */
9034     if (fcf_list.Options) {
9035         mb_fcf_list.options = FCF_LIST_RETURN_ONE;
9036     } else {
9037         mb_fcf_list.options = FCF_LIST_RETURN_ALL;
9038     }

9039     mb_fcf_list.fcf_index = (uint16_t)fcf_list.FcfIndex;
9040     mb_fcf_list.buffer_size = fcf_list.BufSize;

9041     /* Send command */
9042     rval = ql_get_fcf_list_mbx(ha, &mb_fcf_list, (caddr_t)tmp_buf);
9043     if (rval != QL_SUCCESS) {
9044         /* error */
9045         EL(ha, "failed, get_fcf_list_mbx=%xh\n", rval);
9046         kmem_free(tmp_buf, fcf_list.BufSize);
9047         cmd->Status = EXT_STATUS_ERR;
9048         cmd->ResponseLen = 0;
9049         return;
9050     }

9051     /* Copy the response */
9052     if (ql_send_buffer_data((caddr_t)tmp_buf,
9053         (caddr_t)(uintptr_t)cmd->ResponseAdr,
9054         fcf_list.BufSize, mode) != fcf_list.BufSize) {
9055         EL(ha, "failed, ddi_copyout\n");
9056         cmd->Status = EXT_STATUS_COPY_ERR;
9057         cmd->ResponseLen = 0;
9058     } else {
9059         cmd->ResponseLen = mb_fcf_list.buffer_size;
9060         QL_PRINT_9(CE_CONT, "(%d): done\n", ha->instance);
9061     }

9062     kmem_free(tmp_buf, fcf_list.BufSize);

```

new/usr/src/uts/common/io/fibre-channel/fca/qlc/ql_xioctl.c

55

9062 }

unchanged_portion_omitted

```

*****
38588 Thu Oct 23 11:04:56 2014
new/usr/src/uts/common/io/fibre-channel/fca/qlge/qlge_flash.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

213 /*
214 * qlge_load_flash
215 * Write "size" bytes from memory "dp" to flash address "faddr".
216 * faddr = 32bit word flash address.
217 */
218 int
219 qlge_load_flash(qlge_t *qlge, uint8_t *dp, uint32_t len, uint32_t faddr)
220 {
221     int rval = DDI_FAILURE;
222     uint32_t start_block, end_block;
223     uint32_t start_byte, end_byte;
224     uint32_t num;
225     uint32_t sector_size, addr_src, addr_desc;
226     uint8_t *temp;
227     caddr_t bp, bdesc;

229     QL_PRINT(DBG_FLASH, ("%s(%d) entered to write addr %x, %d bytes\n",
230         __func__, qlge->instance, faddr, len));

232     sector_size = qlge->fdesc.block_size;

234     if (faddr > qlge->fdesc.flash_size) {
235         cmn_err(CE_WARN, "%s(%d): invalid flash write address %x",
236             __func__, qlge->instance, faddr);
237         return (DDI_FAILURE);
238     }
239     /* Get semaphore to access Flash Address and Flash Data Registers */
240     if (ql_sem_spinlock(qlge, QL_FLASH_SEM_MASK) != DDI_SUCCESS) {
241         return (DDI_FAILURE);
242     }
243     temp = kmem_zalloc(sector_size, KM_SLEEP);
244     if (temp == NULL) {
245         cmn_err(CE_WARN, "%s(%d): Unable to allocate buffer",
246             __func__, qlge->instance);
247         ql_sem_unlock(qlge, QL_FLASH_SEM_MASK);
248         return (DDI_FAILURE);
249     }

245     (void) ql_unprotect_flash(qlge);

247     get_sector_number(qlge, faddr, &start_block);
248     get_sector_number(qlge, faddr + len - 1, &end_block);

250     QL_PRINT(DBG_FLASH, ("%s(%d) start_block %x, end_block %x\n",
251         __func__, qlge->instance, start_block, end_block));

253     for (num = start_block; num <= end_block; num++) {
254         QL_PRINT(DBG_FLASH,
255             ("%s(%d) sector_size 0x%x, sector read addr %x\n",
256                 __func__, qlge->instance, sector_size, num * sector_size));
257         /* read one whole sector flash data to buffer */
258         rval = qlge_dump_fcode(qlge, (uint8_t *)temp, sector_size,
259             num * sector_size);

261         start_byte = num * sector_size;
262         end_byte = start_byte + sector_size - 1;
263         if (start_byte < faddr)
264             start_byte = faddr;

```

```

265         if (end_byte > (faddr + len))
266             end_byte = (faddr + len - 1);

268         addr_src = start_byte - faddr;
269         addr_desc = start_byte - num * sector_size;
270         bp = (caddr_t)dp + addr_src;
271         bdesc = (caddr_t)temp + addr_desc;
272         bcopy(bp, bdesc, (end_byte - start_byte + 1));

274         /* write the whole sector data to flash */
275         if (ql_erase_and_write_to_flash(qlge, temp, sector_size,
276             num * sector_size) != DDI_SUCCESS)
277             goto out;
278     }
279     rval = DDI_SUCCESS;
280 out:
281     (void) ql_protect_flash(qlge);
282     kmem_free(temp, sector_size);

284     ql_sem_unlock(qlge, QL_FLASH_SEM_MASK);

286     if (rval != DDI_SUCCESS) {
287         cmn_err(CE_WARN, "%s(%d) failed=%xh",
288             __func__, qlge->instance, rval);
289     }

291     return (rval);
292 }
_____unchanged_portion_omitted_____

```

new/usr/src/uts/common/io/iwh/iwh.c

1

153507 Thu Oct 23 11:04:56 2014

new/usr/src/uts/common/io/iwh/iwh.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

unchanged_portion_omitted

1795 /* ARGSUSED */

1796 static ieee80211_node_t *

1797 iwh_node_alloc(ieee80211com_t *ic)

1798 {

1799 iwh_amrr_t *amrr;

1801 amrr = kmem_zalloc(sizeof (iwh_amrr_t), KM_SLEEP);

1802 if (NULL == amrr) {

1803 cmn_err(CE_WARN, "iwh_node_alloc(): "

1804 "failed to allocate memory for amrr structure\n");

1805 return (NULL);

1806 }

1803 iwh_amrr_init(amrr);

1805 return (&amrr->in);

1806 }

unchanged_portion_omitted

new/usr/src/uts/common/io/iwp/iwp.c

1

126658 Thu Oct 23 11:04:56 2014

new/usr/src/uts/common/io/iwp/iwp.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

unchanged_portion_omitted

1732 /* ARGSUSED */

1733 static ieee80211_node_t *

1734 iwp_node_alloc(ieee80211com_t *ic)

1735 {

1736 iwp_amrr_t *amrr;

1738 amrr = kmem_zalloc(sizeof (iwp_amrr_t), KM_SLEEP);

1739 if (NULL == amrr) {

1740 cmn_err(CE_WARN, "iwp_node_alloc(): "

1741 "failed to allocate memory for amrr structure\n");

1742 return (NULL);

1743 }

1740 iwp_amrr_init(amrr);

1742 return (&amrr->in);

1743 }

unchanged_portion_omitted

```

*****
114548 Thu Oct 23 11:04:57 2014
new/usr/src/uts/common/io/mwl/mwl.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

542 static int
543 mwl_alloc_tx_ring(struct mwl_softc *sc, struct mwl_tx_ring *ring,
544                 int count)
545 {
546     struct mwl_txdesc *ds;
547     struct mwl_txbuf *bf;
548     int i, err, datadlen;

550     ring->count = count;
551     ring->queued = 0;
552     ring->cur = ring->next = ring->stat = 0;
553     err = mwl_alloc_dma_mem(sc->sc_dev, &mwl_dma_attr,
554                          count * sizeof (struct mwl_txdesc), &mwl_desc_accattr,
555                          DDI_DMA_CONSISTENT, DDI_DMA_RDWR | DDI_DMA_CONSISTENT,
556                          &ring->txdesc_dma);
557     if (err) {
558         MWL_DBG(MWL_DBG_DMA, "mwl: mwl_alloc_tx_ring(): "
559              "alloc tx ring failed, size %d\n",
560              (uint32_t)(count * sizeof (struct mwl_txdesc)));
561         return (DDI_FAILURE);
562     }

564     MWL_DBG(MWL_DBG_DMA, "mwl: mwl_alloc_tx_ring(): "
565              "dma len = %d\n", (uint32_t)(ring->txdesc_dma.alength));
566     ring->desc = (struct mwl_txdesc *)ring->txdesc_dma.mem_va;
567     ring->physaddr = ring->txdesc_dma.cookie.dmac_address;
568     bzero(ring->desc, count * sizeof (struct mwl_txdesc));

570     datadlen = count * sizeof (struct mwl_txbuf);
571     ring->buf = kmem_zalloc(datadlen, KM_SLEEP);
572     if (ring->buf == NULL) {
573         MWL_DBG(MWL_DBG_DMA, "mwl: mwl_alloc_tx_ring(): "
574              "could not alloc tx ring data buffer\n");
575         return (DDI_FAILURE);
576     }
577     bzero(ring->buf, count * sizeof (struct mwl_txbuf));

574     for (i = 0; i < count; i++) {
575         ds = &ring->desc[i];
576         bf = &ring->buf[i];
577         /* alloc DMA memory */
578         (void) mwl_alloc_dma_mem(sc->sc_dev, &mwl_dma_attr,
579                               sc->sc_dmabuf_size,
580                               &mwl_buf_accattr,
581                               DDI_DMA_STREAMING,
582                               DDI_DMA_WRITE | DDI_DMA_STREAMING,
583                               &bf->txbuf_dma);
584         bf->bf_baddr = bf->txbuf_dma.cookie.dmac_address;
585         bf->bf_mem = (uint8_t *) (bf->txbuf_dma.mem_va);
586         bf->bf_daddr = ring->physaddr + _PTRDIFF(ds, ring->desc);
587         bf->bf_desc = ds;
588     }

590     (void) ddi_dma_sync(ring->txdesc_dma.dma_hdl,
591                       0,
592                       ring->txdesc_dma.alength,
593                       DDI_DMA_SYNC_FORDEV);

```

```

595         return (0);
596     }
_____unchanged_portion_omitted_____

2217 /* ARGSUSED */
2218 static struct ieee80211_node *
2219 mwl_node_alloc(struct ieee80211com *ic)
2220 {
2221     struct mwl_node *mn;

2223     mn = kmem_zalloc(sizeof (struct mwl_node), KM_SLEEP);
2224     if (mn == NULL) {
2225         /* XXX stat+msg */
2226         MWL_DBG(MWL_DBG_MSG, "mwl: mwl_node_alloc(): "
2227              "alloc node failed\n");
2228         return (NULL);
2229     }
2230     return (&mn->mn_node);
2231 }
_____unchanged_portion_omitted_____

```

new/usr/src/uts/common/io/net80211/net80211_crypto_ccmp.c

1

13141 Thu Oct 23 11:04:57 2014

new/usr/src/uts/common/io/net80211/net80211_crypto_ccmp.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

unchanged_portion_omitted

83 /* ARGSUSED */

84 static void *

85 ccmp_attach(struct ieee80211com *ic, struct ieee80211_key *k)

86 {

87 struct ccmp_ctx *ctx;

89 ctx = kmem_zalloc(sizeof (struct ccmp_ctx), KM_SLEEP);

90 if (ctx == NULL)

91 return (NULL);

91 ctx->cc_ic = ic;

92 return (ctx);

93 }

unchanged_portion_omitted

new/usr/src/uts/common/io/net80211/net80211_crypto_tkip.c

1

21087 Thu Oct 23 11:04:57 2014

new/usr/src/uts/common/io/net80211/net80211_crypto_tkip.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

_____unchanged_portion_omitted_____

```
89 static void michael_mic(struct tkip_ctx *, const uint8_t *,
90     mblk_t *, uint_t, size_t, uint8_t[]);
91 static int tkip_encrypt(struct tkip_ctx *, struct ieee80211_key *,
92     mblk_t *, int);
93 static int tkip_decrypt(struct tkip_ctx *, struct ieee80211_key *,
94     mblk_t *, int);

96 extern int rc4_init(crypto_context_t *, const uint8_t *, int);
97 extern int rc4_crypt(crypto_context_t, const uint8_t *, uint8_t *, int);
98 extern int rc4_final(crypto_context_t, uint8_t *, int);
```

```
100 /* ARGSUSED */
101 static void *
102 tkip_attach(struct ieee80211com *ic, struct ieee80211_key *k)
103 {
104     struct tkip_ctx *ctx;

106     ctx = kmem_zalloc(sizeof (struct tkip_ctx), KM_SLEEP);
107     if (ctx == NULL)
108         return (NULL);

108     ctx->tc_ic = ic;
109     return (ctx);
110 }
_____unchanged_portion_omitted_____
```



```

*****
12883 Thu Oct 23 11:04:57 2014
new/usr/src/uts/common/io/ntxn/unm_nic_init.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

253 int
254 pinit_from_rom(struct unm_adapter_s *adapter, int verbose)
255 {
256     int    addr, val, status, i, init_delay = 0, n;
257     struct crb_addr_pair *buf;
258     unsigned long off;
259     unsigned int  offset;

261     status = unm_nic_get_board_info(adapter);
262     if (status)
263         cmn_err(CE_WARN, "%s: pinit_from_rom: Error getting brdinfo\n",
264                unm_nic_driver_name);

266     UNM_CRB_WRITE_LIT_ADAPTER(UNM_ROMUSB_GLB_SW_RESET, 0xffffffff, adapter);

268     if (verbose) {
269         int    val;
270         if (rom_fast_read(adapter, 0x4008, &val) == 0)
271             cmn_err(CE_WARN, "P2 ROM board type: 0x%08x\n", val);
272         else
273             cmn_err(CE_WARN, "Could not read board type\n");
274         if (rom_fast_read(adapter, 0x400c, &val) == 0)
275             cmn_err(CE_WARN, "ROM board num: 0x%08x\n", val);
276         else
277             cmn_err(CE_WARN, "Could not read board number\n");
278         if (rom_fast_read(adapter, 0x4010, &val) == 0)
279             cmn_err(CE_WARN, "ROM chip num: 0x%08x\n", val);
280         else
281             cmn_err(CE_WARN, "Could not read chip number\n");
282     }

284     if (NX_IS_REVISION_P3(adapter->ahw.revision_id)) {
285         if (rom_fast_read(adapter, 0, &n) != 0 ||
286             (unsigned int)n != 0xcafecafe ||
287             rom_fast_read(adapter, 4, &n) != 0) {
288             cmn_err(CE_WARN, "%s: ERROR Reading crb_init area: "
289                    "n: %08x\n", unm_nic_driver_name, n);
290             return (-1);
291         }

293         offset = n & 0xffffU;
294         n = (n >> 16) & 0xffffU;
295     } else {
296         if (rom_fast_read(adapter, 0, &n) != 0 ||
297             !(n & 0x80000000)) {
298             cmn_err(CE_WARN, "%s: ERROR Reading crb_init area: "
299                    "n: %08x\n", unm_nic_driver_name, n);
300             return (-1);
301         }
302         offset = 1;
303         n &= ~0x80000000;
304     }

306     if (n >= 1024) {
307         cmn_err(CE_WARN, "%s: %s:n=0x%x Card flash not initialized\n",
308                unm_nic_driver_name, __FUNCTION__, n);
309         return (-1);
310     }

```

```

312     if (verbose)
313         cmn_err(CE_WARN, "%s: %d CRB init values found in ROM.\n",
314                unm_nic_driver_name, n);

316     buf = kmem_zalloc(n * sizeof (struct crb_addr_pair), KM_SLEEP);
317     if (buf == NULL) {
318         cmn_err(CE_WARN, "%s: pinit_from_rom: Unable to get memory\n",
319                unm_nic_driver_name);
320         return (-1);
321     }

318     for (i = 0; i < n; i++) {
319         if (rom_fast_read(adapter, 8*i + 4*offset, &val) != 0 ||
320             rom_fast_read(adapter, 8*i + 4*offset + 4, &addr) != 0) {
321             kmem_free(buf, n * sizeof (struct crb_addr_pair));
322             return (-1);
323         }

325         buf[i].addr = addr;
326         buf[i].data = val;

328         if (verbose)
329             cmn_err(CE_WARN, "%s: PCI:      0x%08x == 0x%08x\n",
330                    unm_nic_driver_name,
331                    (unsigned int)decode_crb_addr(
332                     (unsigned long)addr), val);
333     }

335     for (i = 0; i < n; i++) {
336         off = decode_crb_addr((unsigned long)buf[i].addr) +
337             UNM_PCI_CRBSPACE;
338         /* skipping cold reboot MAGIC */
339         if (off == UNM_CAM_RAM(0x1fc)) {
340             continue;
341         }

343         if (NX_IS_REVISION_P3(adapter->ahw.revision_id)) {
344             /* do not reset PCI */
345             if (off == (ROMUSB_GLB + 0xbc)) {
346                 continue;
347             }
348             if (off == (ROMUSB_GLB + 0xc8)) /* core clock */
349                 continue;
350             if (off == (ROMUSB_GLB + 0x24)) /* MN clock */
351                 continue;
352             if (off == (ROMUSB_GLB + 0x1c)) /* MS clock */
353                 continue;
354             if (off == (UNM_CRB_PEG_NET_1 + 0x18)) {
355                 buf[i].data = 0x1020;
356             }
357             /* skip the function enable register */
358             if (off == UNM_PCIE_REG(PCIE_SETUP_FUNCTION)) {
359                 continue;
360             }
361             if (off == UNM_PCIE_REG(PCIE_SETUP_FUNCTION2)) {
362                 continue;
363             }

365             if ((off & 0x0ff00000) == UNM_CRB_SMB) {
366                 continue;
367             }

369         }

371         if (off == ADDR_ERROR) {

```

```
372             cmn_err(CE_WARN, "%s: Err: Unknown addr: 0x%08lx\n",
373                   unm_nic_driver_name, buf[i].addr);
374             continue;
375         }
376
377         /* After writing this register, HW needs time for CRB */
378         /* to quiet down (else crb_window returns 0xffffffff) */
379         if (off == UNM_ROMUSB_GLB_SW_RESET) {
380             init_delay = 1;
381
382             if (NX_IS_REVISION_P2(adapter->ahw.revision_id)) {
383                 /* hold xdma in reset also */
384                 buf[i].data = 0x8000ff;
385             }
386         }
387
388         adapter->unm_nic_hw_write_wx(adapter, off, &buf[i].data, 4);
389
390         if (init_delay == 1) {
391             nx_msleep(1000);           /* Sleep 1000 msecs */
392             init_delay = 0;
393         }
394
395         nx_msleep(1);                 /* Sleep 1 msec */
396     }
397
398     kmem_free(buf, n * sizeof (struct crb_addr_pair));
399
400     // disable_peg_cache_all
401     // unreset_net_cache
402     if (NX_IS_REVISION_P2(adapter->ahw.revision_id)) {
403         val = UNM_CRB_READ_VAL_ADAPTER(UNM_ROMUSB_GLB_SW_RESET,
404                                       adapter);
405         UNM_CRB_WRITE_LIT_ADAPTER(UNM_ROMUSB_GLB_SW_RESET,
406                                  (val & 0xfffff0f), adapter);
407     }
408
409     // p2dn replyCount
410     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_D+0xec, 0x1e, adapter);
411     // disable_peg_cache 0
412     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_D+0x4c, 8, adapter);
413     // disable_peg_cache 1
414     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_I+0x4c, 8, adapter);
415
416     // peg_clr_all
417     // peg_clr 0
418     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_0+0x8, 0, adapter);
419     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_0+0xc, 0, adapter);
420     // peg_clr 1
421     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_1+0x8, 0, adapter);
422     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_1+0xc, 0, adapter);
423     // peg_clr 2
424     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_2+0x8, 0, adapter);
425     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_2+0xc, 0, adapter);
426     // peg_clr 3
427     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_3+0x8, 0, adapter);
428     UNM_CRB_WRITE_LIT_ADAPTER(UNM_CRB_PEG_NET_3+0xc, 0, adapter);
429
430     return (0);
431 }
432
433 unchanged_portion_omitted
```

32939 Thu Oct 23 11:04:58 2014

new/usr/src/uts/common/io/pciex/hotplug/pcie_hp.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

unchanged_portion_omitted

```
118 /*
119  * routine to copy in a nvlist from userland
120  */
121 int
122 pcie_copyin_nvlist(char *packed_buf, size_t packed_sz, nvlist_t **nvlp)
123 {
124     int         ret = DDI_SUCCESS;
125     char        *packed;
126     nvlist_t    *dest = NULL;
127
128     if (packed_buf == NULL || packed_sz == 0)
129         return (DDI_EINVAL);
130
131     /* copyin packed nvlist */
132     packed = kmem_alloc(packed_sz, KM_SLEEP);
133     if ((packed = kmem_alloc(packed_sz, KM_SLEEP)) == NULL)
134         return (DDI_ENOMEM);
135
136     if (copyin(packed_buf, packed, packed_sz) != 0) {
137         cmn_err(CE_WARN, "pcie_copyin_nvlist: copyin failed.\n");
138         ret = DDI_FAILURE;
139         goto copyin_cleanup;
140     }
141
142     /* unpack packed nvlist */
143     if ((ret = nvlist_unpack(packed, packed_sz, &dest, KM_SLEEP)) != 0) {
144         cmn_err(CE_WARN, "pcie_copyin_nvlist: nvlist_unpack "
145             "failed with err %d\n", ret);
146         switch (ret) {
147             case EINVAL:
148                 ret = DDI_EINVAL;
149                 goto copyin_cleanup;
150             case ENOMEM:
151                 ret = DDI_ENOMEM;
152                 goto copyin_cleanup;
153             default:
154                 ret = DDI_FAILURE;
155                 goto copyin_cleanup;
156         }
157     }
158     *nvlp = dest;
159     copyin_cleanup:
160     kmem_free(packed, packed_sz);
161     return (ret);
162 }
```

unchanged_portion_omitted

```

*****
86645 Thu Oct 23 11:04:58 2014
new/usr/src/uts/common/io/rtw/rtw.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

786 /*
787  * Returns -1 on failure.
788  */
789 static int
790 rtw_srom_read(struct rtw_regs *regs, uint32_t flags, struct rtw_srom *sr,
791             const char *dvname)
792 {
793     int rc;
794     struct seeprom_descriptor sd;
795     uint8_t ecr;

797     (void) memset(&sd, 0, sizeof (sd));

799     ecr = RTW_READ8(regs, RTW_9346CR);

801     if ((flags & RTW_F_9356SROM) != 0) {
802         RTW_DPRINTF(RTW_DEBUG_ATTACH, "%s: 93c56 SROM\n", dvname);
803         sr->sr_size = 256;
804         sd.sd_chip = C56_66;
805     } else {
806         RTW_DPRINTF(RTW_DEBUG_ATTACH, "%s: 93c46 SROM\n", dvname);
807         sr->sr_size = 128;
808         sd.sd_chip = C46;
809     }

811     ecr &= ~(RTW_9346CR_EEDI | RTW_9346CR_EEDO | RTW_9346CR_EESK |
812            RTW_9346CR_EEM_MASK | RTW_9346CR_EECS);
813     ecr |= RTW_9346CR_EEM_PROGRAM;

815     RTW_WRITE8(regs, RTW_9346CR, ecr);

817     sr->sr_content = kmem_zalloc(sr->sr_size, KM_SLEEP);

819     if (sr->sr_content == NULL) {
820         cmn_err(CE_WARN, "%s: unable to allocate SROM buffer\n",
821             dvname);
822         return (ENOMEM);
823     }

819     (void) memset(sr->sr_content, 0, sr->sr_size);

821     /*
822     * RTL8180 has a single 8-bit register for controlling the
823     * 93cx6 SROM. There is no "ready" bit. The RTL8180
824     * input/output sense is the reverse of read_seeprom's.
825     */
826     sd.sd_handle = regs->r_handle;
827     sd.sd_base = regs->r_base;
828     sd.sd_regsize = 1;
829     sd.sd_control_offset = RTW_9346CR;
830     sd.sd_status_offset = RTW_9346CR;
831     sd.sd_dataout_offset = RTW_9346CR;
832     sd.sd_CK = RTW_9346CR_EESK;
833     sd.sd_CS = RTW_9346CR_EECS;
834     sd.sd_DI = RTW_9346CR_EEDO;
835     sd.sd_DO = RTW_9346CR_EEDI;
836     /*
837     * make read_seeprom enter EEPROM read/write mode

```

```

838     /*
839     sd.sd_MS = ecr;
840     sd.sd_RDY = 0;

842     /*
843     * TBD bus barriers
844     */
845     if (!read_seeprom(&sd, sr->sr_content, 0, sr->sr_size/2)) {
846         cmn_err(CE_WARN, "%s: could not read SROM\n", dvname);
847         kmem_free(sr->sr_content, sr->sr_size);
848         sr->sr_content = NULL;
849         return (-1); /* XXX */
850     }

852     /*
853     * end EEPROM read/write mode
854     */
855     RTW_WRITE8(regs, RTW_9346CR,
856             (ecr & ~RTW_9346CR_EEM_MASK) | RTW_9346CR_EEM_NORMAL);
857     RTW_WBRW(regs, RTW_9346CR, RTW_9346CR);

859     if ((rc = rtw_recall_eeeprom(regs, dvname)) != 0)
860         return (rc);

862 #ifdef SROM_DEBUG
863     {
864         int i;
865         RTW_DPRINTF(RTW_DEBUG_ATTACH,
866             "\n%s: serial ROM:\n\t", dvname);
867         for (i = 0; i < sr->sr_size/2; i++) {
868             RTW_DPRINTF(RTW_DEBUG_ATTACH,
869                 "offset-0x%x: %04x", 2*i, sr->sr_content[i]);
870         }
871     }
872 #endif /* SROM_DEBUG */
873     return (0);
874 }
_____unchanged_portion_omitted_____

```

new/usr/src/uts/common/io/scsi/adapters/mpt_sas/mptsas_impl.c

1

83079 Thu Oct 23 11:04:58 2014

new/usr/src/uts/common/io/scsi/adapters/mpt_sas/mptsas_impl.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

_____unchanged_portion_omitted_____

```
1040 int
1041 mptsas_request_from_pool(mptsas_t *mpt, mptsas_cmd_t **cmd,
1042     struct scsi_pkt **pkt)
1043 {
1044     m_event_struct_t      *ioc_cmd = NULL;
1046     ioc_cmd = kmem_zalloc(M_EVENT_STRUCT_SIZE, KM_SLEEP);
1047     if (ioc_cmd == NULL) {
1048         return (DDI_FAILURE);
1049     }
1047     ioc_cmd->m_event_linkp = NULL;
1048     mptsas_ioc_event_cmdq_add(mpt, ioc_cmd);
1049     *cmd = &(ioc_cmd->m_event_cmd);
1050     *pkt = &(ioc_cmd->m_event_pkt);
1052     return (DDI_SUCCESS);
1053 }
_____unchanged_portion_omitted_____
```

new/usr/src/uts/common/io/scsi/adapters/pmcs/pmcs_attach.c

1

```
*****
84832 Thu Oct 23 11:04:58 2014
new/usr/src/uts/common/io/scsi/adapters/pmcs/pmcs_attach.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____
```

```
1875 static int
1876 pmcs_add_more_chunks(pmcs_hw_t *pwp, unsigned long nsize)
1877 {
1878     pmcs_dmachunk_t *dc;
1879     unsigned long dl;
1880     pmcs_chunk_t *pchunk = NULL;

1882     pwp->cip_dma_attr.dma_attr_align = sizeof (uint32_t);

1884     pchunk = kmem_zalloc(sizeof (pmcs_chunk_t), KM_SLEEP);
1885     if (pchunk == NULL) {
1886         pmcs_prt(pwp, PMCS_PRT_DEBUG, NULL, NULL,
1887             "Not enough memory for DMA chunks");
1888         return (-1);
1889     }

1886     if (pmcs_dma_setup(pwp, &pwp->cip_dma_attr, &pchunk->acc_handle,
1887         &pchunk->dma_handle, nsize, (caddr_t *)&pchunk->addrp,
1888         &pchunk->dma_addr) == B_FALSE) {
1889         pmcs_prt(pwp, PMCS_PRT_DEBUG, NULL, NULL,
1890             "Failed to setup DMA for chunks");
1891         kmem_free(pchunk, sizeof (pmcs_chunk_t));
1892         return (-1);
1893     }

1895     if ((pmcs_check_acc_handle(pchunk->acc_handle) != DDI_SUCCESS) ||
1896         (pmcs_check_dma_handle(pchunk->dma_handle) != DDI_SUCCESS)) {
1897         ddi_fm_service_impact(pwp->dip, DDI_SERVICE_UNAFFECTED);
1898         return (-1);
1899     }

1901     bzero(pchunk->addrp, nsize);
1902     dc = NULL;
1903     for (dl = 0; dl < (nsize / PMCS_SGL_CHUNKSZ); dl++) {
1904         pmcs_dmachunk_t *tmp;
1905         tmp = kmem_alloc(sizeof (pmcs_dmachunk_t), KM_SLEEP);
1906         tmp->nxt = dc;
1907         dc = tmp;
1908     }
1909     mutex_enter(&pwp->dma_lock);
1910     pmcs_idma_chunks(pwp, dc, pchunk, nsize);
1911     pwp->nchunks++;
1912     mutex_exit(&pwp->dma_lock);
1913     return (0);
1914 }
_____unchanged_portion_omitted_____
```

```

*****
31172 Thu Oct 23 11:04:58 2014
new/usr/src/uts/common/io/scsi/targets/ses_safte.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License, Version 1.0 only
6  * (the "License").  You may not use this file except in compliance
7  * with the License.
8  *
9  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
10 * or http://www.opensolaris.org/os/licensing.
11 * See the License for the specific language governing permissions
12 * and limitations under the License.
13 *
14 * When distributing Covered Code, include this CDDL HEADER in each
15 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
16 * If applicable, add the following below this CDDL HEADER, with the
17 * fields enclosed by brackets "[]" replaced with your own identifying
18 * information: Portions Copyright [yyyy] [name of copyright owner]
19 *
20 * CDDL HEADER END
21 */
22 /*
23 * Enclosure Services Devices, SAF-TE Enclosure Routines
24 *
25 * Copyright 2004 Sun Microsystems, Inc.  All rights reserved.
26 * Use is subject to license terms.
27 */

29 #pragma ident      "%Z%%M% %I%      %E% SMI"

29 #include <sys/modctl.h>
30 #include <sys/file.h>
31 #include <sys/scsi/scsi.h>
32 #include <sys/stat.h>
33 #include <sys/scsi/targets/sesio.h>
34 #include <sys/scsi/targets/ses.h>

37 static int set_objstat_sel(ses_softc_t *, ses_objarg *, int);
38 static int wrbuf16(ses_softc_t *, uchar_t, uchar_t, uchar_t, uchar_t, int);
39 static void wrslot_stat(ses_softc_t *, int);
40 static int perf_slotop(ses_softc_t *, uchar_t, uchar_t, int);

42 #define ALL_ENC_STAT \
43     (ENCSTAT_CRITICAL|ENCSTAT_UNRECOV|ENCSTAT_NONCRITICAL|ENCSTAT_INFO)

45 #define SCRATCH 64
46 #define NPSEUDO_THERM 1
47 #define NPSEUDO_ALARM 1
48 struct scfg {
49     /*
50      * Cached Configuration
51      */
52     uchar_t Nfans;          /* Number of Fans */
53     uchar_t Npwr;          /* Number of Power Supplies */
54     uchar_t Nslots;        /* Number of Device Slots */
55     uchar_t DoorLock;      /* Door Lock Installed */
56     uchar_t Ntherm;        /* Number of Temperature Sensors */
57     uchar_t Nspkr;        /* Number of Speakers */
58     uchar_t Nalarm;        /* Number of Alarms (at least one) */

```

```

59     /*
60      * Cached Flag Bytes for Global Status
61      */
62     uchar_t flag1;
63     uchar_t flag2;
64     /*
65      * What object index ID is where various slots start.
66      */
67     uchar_t pwoff;
68     uchar_t slotoff;
69 #define ALARM_OFFSET(cc)          (cc)->slotoff - 1
70 };
71 #define FLG1_ALARM          0x1
72 #define FLG1_GLOBFAIL      0x2
73 #define FLG1_GLOBWARN      0x4
74 #define FLG1_ENCPWROFF      0x8
75 #define FLG1_ENCFANFAIL    0x10
76 #define FLG1_ENCPWRFAIL    0x20
77 #define FLG1_ENCDRVFAIL    0x40
78 #define FLG1_ENCDRVWARN    0x80

80 #define FLG2_LOCKDOOR      0x4
81 #define SAFTE_PRIVATE      sizeof (struct scfg)

83 #if !defined(lint)
84 _NOTE(MUTEX_PROTECTS_DATA(scsi_device::sd_mutex, scfg))
85 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:Nfans))
86 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:Npwr))
87 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:Nslots))
88 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:DoorLock))
89 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:Ntherm))
90 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:Nspkr))
91 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:Nalarm))
92 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:flag1))
93 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:flag2))
94 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:pwoff))
95 _NOTE(DATA_READABLE_WITHOUT_LOCK(scfg:slotoff))
96 #endif

98 static int
99 safte_getconfig(ses_softc_t *ssc)
100 {
101     struct scfg *cfg;
102     int err;
103     Uscmd local, *lp = &local;
104     char rqbuf[SENSE_LENGTH], *sdata;
105     static char cdb[CDB_GROUP1] =
106         { SCMD_READ_BUFFER, 1, SAFTE_RD_RDCFG, 0, 0, 0, 0, 0, 0, SCRATCH, 0 };

108     cfg = ssc->ses_private;
109     if (cfg == NULL)
110         return (ENXIO);

112     sdata = kmem_alloc(SCRATCH, KM_SLEEP);
113     if (sdata == NULL)
114         return (ENOMEM);

114     lp->uscsi_flags = USCSI_READ|USCSI_RQENABLE;
115     lp->uscsi_timeout = ses_io_time;
116     lp->uscsi_cdb = cdb;
117     lp->uscsi_bufaddr = sdata;
118     lp->uscsi_buflen = SCRATCH;
119     lp->uscsi_cdblen = sizeof (cdb);
120     lp->uscsi_rqbuf = rqbuf;
121     lp->uscsi_rqlen = sizeof (rqbuf);

```

```

123     err = ses_runcmd(ssc, lp);
124     if (err) {
125         kmem_free(sdata, SCRATCH);
126         return (err);
127     }
128
129     if ((lp->uscsci_buflen - lp->uscsci_resid) < 6) {
130         SES_LOG(ssc, CE_NOTE, "Too little data (%ld) for configuration",
131             lp->uscsci_buflen - lp->uscsci_resid);
132         kmem_free(sdata, SCRATCH);
133         return (EIO);
134     }
135     SES_LOG(ssc, SES_CE_DEBUG1, "Nfans %d Npwr %d Nslots %d Lck %d Ntherm "
136         "%d Nspkrs %d", sdata[0], sdata[1], sdata[2], sdata[3], sdata[4],
137         sdata[5]);
138
139     mutex_enter(&ssc->ses_devp->sd_mutex);
140     cfg->Nfans = sdata[0];
141     cfg->Npwr = sdata[1];
142     cfg->Nslots = sdata[2];
143     cfg->DoorLock = sdata[3];
144     cfg->Ntherm = sdata[4];
145     cfg->Nspkrs = sdata[5];
146     cfg->Nalarm = NPSEUDO_ALARM;
147     mutex_exit(&ssc->ses_devp->sd_mutex);
148     kmem_free(sdata, SCRATCH);
149     return (0);
150 }
151
152 int
153 safte_softc_init(ses_softc_t *ssc, int doinit)
154 {
155     int r, i;
156     struct scfg *cc;
157
158     if (doinit == 0) {
159         mutex_enter(&ssc->ses_devp->sd_mutex);
160         if (ssc->ses_nobjects) {
161             if (ssc->ses_objmap) {
162                 kmem_free(ssc->ses_objmap,
163                     ssc->ses_nobjects * sizeof (encobj));
164                 ssc->ses_objmap = NULL;
165             }
166             ssc->ses_nobjects = 0;
167         }
168         if (ssc->ses_private) {
169             kmem_free(ssc->ses_private, SAFTE_PRIVATE);
170             ssc->ses_private = NULL;
171         }
172         mutex_exit(&ssc->ses_devp->sd_mutex);
173         return (0);
174     }
175
176     mutex_enter(&ssc->ses_devp->sd_mutex);
177     if (ssc->ses_private == NULL) {
178         ssc->ses_private = kmem_zalloc(SAFTE_PRIVATE, KM_SLEEP);
179         if (ssc->ses_private == NULL) {
180             mutex_exit(&ssc->ses_devp->sd_mutex);
181             return (ENOMEM);
182         }
183     }
184
185     ssc->ses_nobjects = 0;
186     ssc->ses_encstat = 0;
187     mutex_exit(&ssc->ses_devp->sd_mutex);

```

```

185     if ((r = safte_getconfig(ssc)) != 0) {
186         return (r);
187     }
188
189     /*
190     * The number of objects here, as well as that reported by the
191     * READ_BUFFER/GET_CONFIG call, are the over-temperature flags (15)
192     * that get reported during READ_BUFFER/READ_ENC_STATUS.
193     */
194     mutex_enter(&ssc->ses_devp->sd_mutex);
195     cc = ssc->ses_private;
196     ssc->ses_nobjects = cc->Nfans + cc->Npwr + cc->Nslots + cc->DoorLock +
197         cc->Ntherm + cc->Nspkrs + NPSEUDO_THERM + NPSEUDO_ALARM;
198     ssc->ses_objmap = (encobj *)
199         kmem_zalloc(ssc->ses_nobjects * sizeof (encobj), KM_SLEEP);
200     mutex_exit(&ssc->ses_devp->sd_mutex);
201     if (ssc->ses_objmap == NULL)
202         return (ENOMEM);
203
204     r = 0;
205     /*
206     * Note that this is all arranged for the convenience
207     * in later fetches of status.
208     */
209     mutex_enter(&ssc->ses_devp->sd_mutex);
210     for (i = 0; i < cc->Nfans; i++)
211         ssc->ses_objmap[r++].enttype = SESTYP_FAN;
212     cc->pwroff = (uchar_t)r;
213     for (i = 0; i < cc->Npwr; i++)
214         ssc->ses_objmap[r++].enttype = SESTYP_POWER;
215     for (i = 0; i < cc->DoorLock; i++)
216         ssc->ses_objmap[r++].enttype = SESTYP_DOORLOCK;
217     for (i = 0; i < cc->Nspkrs; i++)
218         ssc->ses_objmap[r++].enttype = SESTYP_ALARM;
219     for (i = 0; i < cc->Ntherm; i++)
220         ssc->ses_objmap[r++].enttype = SESTYP_THERM;
221     for (i = 0; i < NPSEUDO_THERM; i++)
222         ssc->ses_objmap[r++].enttype = SESTYP_THERM;
223     cc->slotoff = (uchar_t)r;
224     for (i = 0; i < cc->Nslots; i++)
225         ssc->ses_objmap[r++].enttype = SESTYP_DEVICE;
226     mutex_exit(&ssc->ses_devp->sd_mutex);
227     return (0);
228 }

```

_____unchanged_portion_omitted_____


```

*****
140795 Thu Oct 23 11:04:59 2014
new/usr/src/uts/common/io/usb/hcd/uhci/uhciutil.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

3101 /*
3102  * uhci_alloc_memory_for_tds:
3103  *   - Allocates memory for the isoc/bulk td pools.
3104  */
3105 static int
3106 uhci_alloc_memory_for_tds(
3107     uhci_state_t      *uhcip,
3108     uint_t            num_tds,
3109     uhci_bulk_isoc_xfer_t *info)
3110 {
3111     int                result, i, j, err;
3112     size_t             real_length;
3113     uint_t             ccount, num;
3114     ddi_device_acc_attr_t dev_attr;
3115     uhci_bulk_isoc_td_pool_t *td_pool_ptr1, *td_pool_ptr2;

3117     USB_DPRINTF_L4(PRINT_MASK_ATT, uhcip->uhci_log_hdl,
3118         "uhci_alloc_memory_for_tds: num_tds: 0x%x info: 0x%p "
3119         "num_pools: %u", num_tds, (void *)info, info->num_pools);

3121     /* The host controller will be little endian */
3122     dev_attr.devacc_attr_version = DDI_DEVICE_ATTR_V0;
3123     dev_attr.devacc_attr_endian_flags = DDI_STRUCTURE_LE_ACC;
3124     dev_attr.devacc_attr_dataorder = DDI_STRICTORDER_ACC;

3126     /* Allocate the TD pool structures */
3127     info->td_pools = kmem_zalloc((sizeof (uhci_bulk_isoc_td_pool_t) *
3128         info->num_pools), KM_SLEEP);
3127     if ((info->td_pools = kmem_zalloc(
3128         (sizeof (uhci_bulk_isoc_td_pool_t) * info->num_pools),
3129         KM_SLEEP)) == NULL) {
3130         USB_DPRINTF_L2(PRINT_MASK_ATT, uhcip->uhci_log_hdl,
3131             "uhci_alloc_memory_for_tds: alloc td_pools failed");
3133         return (USB_FAILURE);
3134     }

3130     for (i = 0; i < info->num_pools; i++) {
3131         if (info->num_pools == 1) {
3132             num = num_tds;
3133         } else if (i < (info->num_pools - 1)) {
3134             num = UHCI_MAX_TD_NUM_PER_POOL;
3135         } else {
3136             num = (num_tds % UHCI_MAX_TD_NUM_PER_POOL);
3137         }

3139         td_pool_ptr1 = &info->td_pools[i];

3141         /* Allocate the bulk TD pool DMA handle */
3142         if (ddi_dma_alloc_handle(uhcip->uhci_dip,
3143             &uhcip->uhci_dma_attr, DDI_DMA_SLEEP, 0,
3144             &td_pool_ptr1->dma_handle) != DDI_SUCCESS) {

3146             for (j = 0; j < i; j++) {
3147                 td_pool_ptr2 = &info->td_pools[j];
3148                 result = ddi_dma_unbind_handle(
3149                     td_pool_ptr2->dma_handle);

```

```

3150         ASSERT(result == DDI_SUCCESS);
3151         ddi_dma_mem_free(&td_pool_ptr2->mem_handle);
3152         ddi_dma_free_handle(&td_pool_ptr2->dma_handle);
3153     }

3155     kmem_free(info->td_pools,
3156         (sizeof (uhci_bulk_isoc_td_pool_t) *
3157             info->num_pools));

3159     return (USB_FAILURE);
3160 }

3162     /* Allocate the memory for the bulk TD pool */
3163     if (ddi_dma_mem_alloc(td_pool_ptr1->dma_handle,
3164         num * sizeof (uhci_td_t), &dev_attr,
3165         DDI_DMA_CONSISTENT, DDI_DMA_SLEEP, 0,
3166         &td_pool_ptr1->pool_addr, &real_length,
3167         &td_pool_ptr1->mem_handle) != DDI_SUCCESS) {

3169         ddi_dma_free_handle(&td_pool_ptr1->dma_handle);

3171         for (j = 0; j < i; j++) {
3172             td_pool_ptr2 = &info->td_pools[j];
3173             result = ddi_dma_unbind_handle(
3174                 td_pool_ptr2->dma_handle);
3175             ASSERT(result == DDI_SUCCESS);
3176             ddi_dma_mem_free(&td_pool_ptr2->mem_handle);
3177             ddi_dma_free_handle(&td_pool_ptr2->dma_handle);
3178         }

3180         kmem_free(info->td_pools,
3181             (sizeof (uhci_bulk_isoc_td_pool_t) *
3182                 info->num_pools));

3184         return (USB_FAILURE);
3185     }

3187     /* Map the bulk TD pool into the I/O address space */
3188     result = ddi_dma_addr_bind_handle(td_pool_ptr1->dma_handle,
3189         NULL, (caddr_t)td_pool_ptr1->pool_addr, real_length,
3190         DDI_DMA_RDWR | DDI_DMA_CONSISTENT, DDI_DMA_SLEEP, NULL,
3191         &td_pool_ptr1->cookie, &ccount);

3193     /* Process the result */
3194     err = USB_SUCCESS;

3196     if (result != DDI_DMA_MAPPED) {
3197         USB_DPRINTF_L2(PRINT_MASK_ATT, uhcip->uhci_log_hdl,
3198             "uhci_allocate_memory_for_tds: Result = %d",
3199             result);
3200         uhci_decode_ddi_dma_addr_bind_handle_result(uhcip,
3201             result);

3203         err = USB_FAILURE;
3204     }

3206     if ((result == DDI_DMA_MAPPED) && (ccount != 1)) {
3207         /* The cookie count should be 1 */
3208         USB_DPRINTF_L2(PRINT_MASK_ATT,
3209             uhcip->uhci_log_hdl,
3210             "uhci_allocate_memory_for_tds: "
3211             "More than 1 cookie");

3213         result = ddi_dma_unbind_handle(
3214             td_pool_ptr1->dma_handle);
3215         ASSERT(result == DDI_SUCCESS);

```

```
3217         err = USB_FAILURE;
3218     }
3219
3220     if (err == USB_FAILURE) {
3221
3222         ddi_dma_mem_free(&td_pool_ptr1->mem_handle);
3223         ddi_dma_free_handle(&td_pool_ptr1->dma_handle);
3224
3225         for (j = 0; j < i; j++) {
3226             td_pool_ptr2 = &info->td_pools[j];
3227             result = ddi_dma_unbind_handle(
3228                 td_pool_ptr2->dma_handle);
3229             ASSERT(result == DDI_SUCCESS);
3230             ddi_dma_mem_free(&td_pool_ptr2->mem_handle);
3231             ddi_dma_free_handle(&td_pool_ptr2->dma_handle);
3232         }
3233
3234         kmem_free(info->td_pools,
3235             (sizeof (uhci_bulk_isoc_td_pool_t) *
3236             info->num_pools));
3237
3238         return (USB_FAILURE);
3239     }
3240
3241     bzero((void *)td_pool_ptr1->pool_addr,
3242         num * sizeof (uhci_td_t));
3243     td_pool_ptr1->num_tds = (ushort_t)num;
3244 }
3245
3246     return (USB_SUCCESS);
3247 }
_____unchanged_portion_omitted_____
```

```

*****
72910 Thu Oct 23 11:04:59 2014
new/usr/src/uts/common/io/usb/usba/whcdi.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

1418 /* Acknowledge WUSB Device Disconnect notification, refer to WUSB 1.0/7.6.2 */
1419 int
1420 wusb_hc_ack_disconn(wusb_hc_data_t *hc_data, uint8_t addr)
1421 {
1422     wusb_ie_dev_disconnect_t    *disconn_ie;
1423     uint8_t                      iehdl;
1424     int                          rval;

1426     ASSERT(mutex_owned(&hc_data->hc_mutex));

1428     /*
1429      * the scheme ensures each time only one device addr
1430      * is set each time
1431      */
1432     disconn_ie = kmem_zalloc(sizeof (wusb_ie_dev_disconnect_t), KM_SLEEP);
1433     if (!disconn_ie) {
1434         return (USB_NO_RESOURCES);
1435     }

1434     disconn_ie->bIEIdentifier = WUSB_IE_DEV_DISCONNECT;
1435     disconn_ie->bDeviceAddress[0] = addr;
1436     /* padding, no active wusb device addr will be 1 */
1437     disconn_ie->bDeviceAddress[1] = 1;
1438     disconn_ie->bLength = 4;

1440     rval = wusb_hc_get_iehdl(hc_data, (wusb_ie_header_t *)disconn_ie,
1441                             &iehdl);
1442     if (rval != USB_SUCCESS) {
1443         USB_DPRINTF_L2(DPRINT_MASK_WHCDI, whcdi_log_handle,
1444                       "wusb_hc_ack_disconn: get ie handle fails");
1445         kmem_free(disconn_ie, sizeof (wusb_ie_dev_disconnect_t));

1447         return (rval);
1448     }

1450     rval = wusb_hc_add_mmc_ie(hc_data, 0, 0, iehdl,
1451                             disconn_ie->bLength, (uint8_t *)disconn_ie);
1452     if (rval != USB_SUCCESS) {
1453         USB_DPRINTF_L2(DPRINT_MASK_WHCDI, whcdi_log_handle,
1454                       "wusb_hc_ack_disconn: add dev disconnect ie fails");
1455         wusb_hc_free_iehdl(hc_data, iehdl);
1456         kmem_free(disconn_ie, sizeof (wusb_ie_dev_disconnect_t));

1458         return (rval);
1459     }

1461     mutex_exit(&hc_data->hc_mutex);
1462     /*
1463      * WUSB 1.0/7.5.4 requires the IE to be transmitted at least
1464      * 100ms before ceasing, wait for 150ms here
1465      */
1466     delay(drv_usectohz(150000));
1467     mutex_enter(&hc_data->hc_mutex);

1469     /* cease transmitting the IE */
1470     (void) wusb_hc_remove_mmc_ie(hc_data, (uint8_t)iehdl);
1471     wusb_hc_free_iehdl(hc_data, iehdl);
1472     kmem_free(disconn_ie, sizeof (wusb_ie_dev_disconnect_t));

```

```

1474         return (USB_SUCCESS);
1475     }
_____unchanged_portion_omitted_____

```

69395 Thu Oct 23 11:04:59 2014

new/usr/src/uts/common/io/xge/drv/xgell.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

_____ unchanged_portion_omitted _____

```

2492 static int
2493 xgell_stats_get(queue_t *q, mblk_t *mp, caddr_t cp, cred_t *credp)
2494 {
2495     xgelldev_t *lldev = (xgelldev_t *)cp;
2496     xge_hal_status_e status;
2497     int count = 0, rethesize;
2498     char *buf;

2500     buf = kmem_alloc(XGELL_STATS_BUFSIZE, KM_SLEEP);
2501     if (buf == NULL) {
2502         return (ENOSPC);
2503     }

2502     status = xge_hal_aux_stats_tmac_read(lldev->devh, XGELL_STATS_BUFSIZE,
2503     buf, &rethesize);
2504     if (status != XGE_HAL_OK) {
2505         kmem_free(buf, XGELL_STATS_BUFSIZE);
2506         xge_debug_ll(XGE_ERR, "tmac_read(): status %d", status);
2507         return (EINVAL);
2508     }
2509     count += rethesize;

2511     status = xge_hal_aux_stats_rmac_read(lldev->devh,
2512     XGELL_STATS_BUFSIZE - count,
2513     buf+count, &rethesize);
2514     if (status != XGE_HAL_OK) {
2515         kmem_free(buf, XGELL_STATS_BUFSIZE);
2516         xge_debug_ll(XGE_ERR, "rmac_read(): status %d", status);
2517         return (EINVAL);
2518     }
2519     count += rethesize;

2521     status = xge_hal_aux_stats_pci_read(lldev->devh,
2522     XGELL_STATS_BUFSIZE - count, buf + count, &rethesize);
2523     if (status != XGE_HAL_OK) {
2524         kmem_free(buf, XGELL_STATS_BUFSIZE);
2525         xge_debug_ll(XGE_ERR, "pci_read(): status %d", status);
2526         return (EINVAL);
2527     }
2528     count += rethesize;

2530     status = xge_hal_aux_stats_sw_dev_read(lldev->devh,
2531     XGELL_STATS_BUFSIZE - count, buf + count, &rethesize);
2532     if (status != XGE_HAL_OK) {
2533         kmem_free(buf, XGELL_STATS_BUFSIZE);
2534         xge_debug_ll(XGE_ERR, "sw_dev_read(): status %d", status);
2535         return (EINVAL);
2536     }
2537     count += rethesize;

2539     status = xge_hal_aux_stats_hal_read(lldev->devh,
2540     XGELL_STATS_BUFSIZE - count, buf + count, &rethesize);
2541     if (status != XGE_HAL_OK) {
2542         kmem_free(buf, XGELL_STATS_BUFSIZE);
2543         xge_debug_ll(XGE_ERR, "pci_read(): status %d", status);
2544         return (EINVAL);
2545     }
2546     count += rethesize;

```

```

2548     *(buf + count - 1) = '\0'; /* remove last '\n' */
2549     (void) mi_mpprintf(mp, "%s", buf);
2550     kmem_free(buf, XGELL_STATS_BUFSIZE);

2552     return (0);
2553 }

2555 static int
2556 xgell_pciconf_get(queue_t *q, mblk_t *mp, caddr_t cp, cred_t *credp)
2557 {
2558     xgelldev_t *lldev = (xgelldev_t *)cp;
2559     xge_hal_status_e status;
2560     int rethesize;
2561     char *buf;

2563     buf = kmem_alloc(XGELL_PCICONF_BUFSIZE, KM_SLEEP);
2564     if (buf == NULL) {
2565         return (ENOSPC);
2566     }
2567     status = xge_hal_aux_pci_config_read(lldev->devh, XGELL_PCICONF_BUFSIZE,
2568     buf, &rethesize);
2569     if (status != XGE_HAL_OK) {
2570         kmem_free(buf, XGELL_PCICONF_BUFSIZE);
2571         xge_debug_ll(XGE_ERR, "pci_config_read(): status %d", status);
2572         return (EINVAL);
2573     }
2574     *(buf + rethesize - 1) = '\0'; /* remove last '\n' */
2575     (void) mi_mpprintf(mp, "%s", buf);
2576     kmem_free(buf, XGELL_PCICONF_BUFSIZE);

2577     return (0);
2578 }

2578 static int
2579 xgell_about_get(queue_t *q, mblk_t *mp, caddr_t cp, cred_t *credp)
2580 {
2581     xgelldev_t *lldev = (xgelldev_t *)cp;
2582     xge_hal_status_e status;
2583     int rethesize;
2584     char *buf;

2586     buf = kmem_alloc(XGELL_ABOUT_BUFSIZE, KM_SLEEP);
2587     if (buf == NULL) {
2588         return (ENOSPC);
2589     }
2590     status = xge_hal_aux_about_read(lldev->devh, XGELL_ABOUT_BUFSIZE,
2591     buf, &rethesize);
2592     if (status != XGE_HAL_OK) {
2593         kmem_free(buf, XGELL_ABOUT_BUFSIZE);
2594         xge_debug_ll(XGE_ERR, "about_read(): status %d", status);
2595         return (EINVAL);
2596     }
2597     *(buf + rethesize - 1) = '\0'; /* remove last '\n' */
2598     (void) mi_mpprintf(mp, "%s", buf);
2599     kmem_free(buf, XGELL_ABOUT_BUFSIZE);

2598     return (0);
2599 }

2601 static unsigned long bar0_offset = 0x110; /* adapter_control */

2603 static int
2604 xgell_bar0_get(queue_t *q, mblk_t *mp, caddr_t cp, cred_t *credp)
2605 {
2606     xgelldev_t *lldev = (xgelldev_t *)cp;

```

```

2607     xge_hal_status_e status;
2608     int  retsize;
2609     char *buf;

2611     buf = kmem_alloc(XGELL_IOCTL_BUFSIZE, KM_SLEEP);
2612     if (buf == NULL) {
2621         return (ENOSPC);
2622     }
2612     status = xge_hal_aux_bar0_read(lldev->devh, bar0_offset,
2613     XGELL_IOCTL_BUFSIZE, buf, &retsize);
2614     if (status != XGE_HAL_OK) {
2615         kmem_free(buf, XGELL_IOCTL_BUFSIZE);
2616         xge_debug_ll(XGE_ERR, "bar0_read(): status %d", status);
2617         return (EINVAL);
2618     }
2619     *(buf + retsize - 1) = '\0'; /* remove last '\n' */
2620     (void) mi_mprintf(mp, "%s", buf);
2621     kmem_free(buf, XGELL_IOCTL_BUFSIZE);

2623     return (0);
2624 }
_____ unchanged_portion_omitted _____

2648 static int
2649 xgell_debug_level_get(queue_t *q, mblk_t *mp, caddr_t cp, cred_t *credp)
2650 {
2651     char *buf;

2653     buf = kmem_alloc(XGELL_IOCTL_BUFSIZE, KM_SLEEP);
2654     if (buf == NULL) {
2666         return (ENOSPC);
2667     }
2654     (void) mi_mprintf(mp, "debug_level %d", xge_hal_driver_debug_level());
2655     kmem_free(buf, XGELL_IOCTL_BUFSIZE);

2657     return (0);
2658 }
_____ unchanged_portion_omitted _____

2677 static int
2678 xgell_debug_module_mask_get(queue_t *q, mblk_t *mp, caddr_t cp, cred_t *credp)
2679 {
2680     char *buf;

2682     buf = kmem_alloc(XGELL_IOCTL_BUFSIZE, KM_SLEEP);
2683     if (buf == NULL) {
2699         return (ENOSPC);
2700     }
2683     (void) mi_mprintf(mp, "debug_module_mask 0x%08x",
2684     xge_hal_driver_debug_module_mask());
2685     kmem_free(buf, XGELL_IOCTL_BUFSIZE);

2687     return (0);
2688 }
_____ unchanged_portion_omitted _____

2712 static int
2713 xgell_devconfig_get(queue_t *q, mblk_t *mp, caddr_t cp, cred_t *credp)
2714 {
2715     xgelldev_t *lldev = (xgelldev_t *) (void *) cp;
2716     xge_hal_status_e status;
2717     int  retsize;
2718     char *buf;

2720     buf = kmem_alloc(XGELL_DEVCONF_BUFSIZE, KM_SLEEP);
2739     if (buf == NULL) {

```

```

2740         return (ENOSPC);
2741     }
2721     status = xge_hal_aux_device_config_read(lldev->devh,
2722     XGELL_DEVCONF_BUFSIZE, buf, &retsize);
2723     if (status != XGE_HAL_OK) {
2724         kmem_free(buf, XGELL_DEVCONF_BUFSIZE);
2725         xge_debug_ll(XGE_ERR, "device_config_read(): status %d",
2726         status);
2727         return (EINVAL);
2728     }
2729     *(buf + retsize - 1) = '\0'; /* remove last '\n' */
2730     (void) mi_mprintf(mp, "%s", buf);
2731     kmem_free(buf, XGELL_DEVCONF_BUFSIZE);

2733     return (0);
2734 }
_____ unchanged_portion_omitted _____

```

```

*****
93566 Thu Oct 23 11:04:59 2014
new/usr/src/uts/common/io/yge/yge.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

989 static int
990 yge_add_intr(yge_dev_t *dev, int intr_type)
991 {
992     dev_info_t      *dip;
993     int              count;
994     int              actual;
995     int              rv;
996     int              i, j;

998     dip = dev->d_dip;

1000     rv = ddi_intr_get_nintrs(dip, intr_type, &count);
1001     if ((rv != DDI_SUCCESS) || (count == 0)) {
1002         yge_error(dev, NULL,
1003             "ddi_intr_get_nintrs failed, rv %d, count %d", rv, count);
1004     }
1005     return (DDI_FAILURE);

1007     /*
1008     * Allocate the interrupt. Note that we only bother with a single
1009     * interrupt. One could argue that for MSI devices with dual ports,
1010     * it would be nice to have a separate interrupt per port. But right
1011     * now I don't know how to configure that, so we'll just settle for
1012     * a single interrupt.
1013     */
1014     dev->d_intrcnt = 1;

1016     dev->d_intrsize = count * sizeof (ddi_intr_handle_t);
1017     dev->d_intrh = kmem_zalloc(dev->d_intrsize, KM_SLEEP);
1018     if (dev->d_intrh == NULL) {
1019         yge_error(dev, NULL, "Unable to allocate interrupt handle");
1020         return (DDI_FAILURE);
1021     }

1019     rv = ddi_intr_alloc(dip, dev->d_intrh, intr_type, 0, dev->d_intrcnt,
1020         &actual, DDI_INTR_ALLOC_STRICT);
1021     if ((rv != DDI_SUCCESS) || (actual == 0)) {
1022         yge_error(dev, NULL,
1023             "Unable to allocate interrupt, %d, count %d",
1024             rv, actual);
1025         kmem_free(dev->d_intrh, dev->d_intrsize);
1026         return (DDI_FAILURE);
1027     }

1029     if ((rv = ddi_intr_get_pri(dev->d_intrh[0], &dev->d_intrpri)) !=
1030         DDI_SUCCESS) {
1031         for (i = 0; i < dev->d_intrcnt; i++)
1032             (void) ddi_intr_free(dev->d_intrh[i]);
1033         yge_error(dev, NULL,
1034             "Unable to get interrupt priority, %d", rv);
1035         kmem_free(dev->d_intrh, dev->d_intrsize);
1036         return (DDI_FAILURE);
1037     }

1039     if ((rv = ddi_intr_get_cap(dev->d_intrh[0], &dev->d_intrcap)) !=
1040         DDI_SUCCESS) {
1041         yge_error(dev, NULL,
1042             "Unable to get interrupt capabilities, %d", rv);

```

```

1043         for (i = 0; i < dev->d_intrcnt; i++)
1044             (void) ddi_intr_free(dev->d_intrh[i]);
1045         kmem_free(dev->d_intrh, dev->d_intrsize);
1046         return (DDI_FAILURE);
1047     }

1049     /* register interrupt handler to kernel */
1050     for (i = 0; i < dev->d_intrcnt; i++) {
1051         if ((rv = ddi_intr_add_handler(dev->d_intrh[i], yge_intr,
1052             dev, NULL)) != DDI_SUCCESS) {
1053             yge_error(dev, NULL,
1054                 "Unable to add interrupt handler, %d", rv);
1055             for (j = 0; j < i; j++)
1056                 (void) ddi_intr_remove_handler(dev->d_intrh[j]);
1057             for (i = 0; i < dev->d_intrcnt; i++)
1058                 (void) ddi_intr_free(dev->d_intrh[i]);
1059             kmem_free(dev->d_intrh, dev->d_intrsize);
1060             return (DDI_FAILURE);
1061         }
1062     }

1064     mutex_init(&dev->d_rxlock, NULL, MUTEX_DRIVER,
1065         DDI_INTR_PRI(dev->d_intrpri));
1066     mutex_init(&dev->d_txlock, NULL, MUTEX_DRIVER,
1067         DDI_INTR_PRI(dev->d_intrpri));
1068     mutex_init(&dev->d_phylock, NULL, MUTEX_DRIVER,
1069         DDI_INTR_PRI(dev->d_intrpri));
1070     mutex_init(&dev->d_task_mtx, NULL, MUTEX_DRIVER,
1071         DDI_INTR_PRI(dev->d_intrpri));

1073     return (DDI_SUCCESS);
1074 }
_____unchanged_portion_omitted_____

```

```

*****
248730 Thu Oct 23 11:05:00 2014
new/usr/src/uts/common/os/sunddi.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

5510 /*
5511  * failing to remove a minor node is not of interest
5512  * therefore we do not generate an error message
5513  */
5514 static int
5515 i_log_devfs_minor_remove(dev_info_t *dip, char *minor_name)
5516 {
5517     char *pathname, *class_name;
5518     sysevent_t *ev;
5519     sysevent_id_t eid;
5520     sysevent_value_t se_val;
5521     sysevent_attr_list_t *ev_attr_list = NULL;

5523     /*
5524     * only log ddi_remove_minor_node() calls outside the scope
5525     * of attach/detach reconfigurations and when the dip is
5526     * still initialized.
5527     */
5528     if (DEVI_IS_ATTACHING(dip) || DEVI_IS_DETACHING(dip) ||
5529         (i_ddi_node_state(dip) < DS_INITIALIZED)) {
5530         return (DDI_SUCCESS);
5531     }

5533     i_ddi_di_cache_invalidate();

5535     ev = sysevent_alloc(EC_DEVFS, ESC_DEVFS_MINOR_REMOVE, EP_DDI, SE_SLEEP);
5536     if (ev == NULL) {
5537         return (DDI_SUCCESS);
5538     }

5540     pathname = kmem_alloc(MAXPATHLEN, KM_SLEEP);
5541     if (pathname == NULL) {
5542         sysevent_free(ev);
5543         return (DDI_SUCCESS);
5544     }

5542     (void) ddi_pathname(dip, pathname);
5543     ASSERT(strlen(pathname));
5544     se_val.value_type = SE_DATA_TYPE_STRING;
5545     se_val.value.sv_string = pathname;
5546     if (sysevent_add_attr(&ev_attr_list, DEVFS_PATHNAME,
5547         &se_val, SE_SLEEP) != 0) {
5548         kmem_free(pathname, MAXPATHLEN);
5549         sysevent_free(ev);
5550         return (DDI_SUCCESS);
5551     }

5553     kmem_free(pathname, MAXPATHLEN);

5555     /*
5556     * allow for NULL minor names
5557     */
5558     if (minor_name != NULL) {
5559         se_val.value.sv_string = minor_name;
5560         if (sysevent_add_attr(&ev_attr_list, DEVFS_MINOR_NAME,
5561             &se_val, SE_SLEEP) != 0) {
5562             sysevent_free_attr(ev_attr_list);
5563             goto fail;

```

```

5564     }
5565     }

5567     if ((class_name = i_ddi_devi_class(dip)) != NULL) {
5568         /* add the device class, driver name and instance attributes */

5570         se_val.value_type = SE_DATA_TYPE_STRING;
5571         se_val.value.sv_string = class_name;
5572         if (sysevent_add_attr(&ev_attr_list,
5573             DEVFS_DEVI_CLASS, &se_val, SE_SLEEP) != 0) {
5574             sysevent_free_attr(ev_attr_list);
5575             goto fail;
5576         }

5578         se_val.value_type = SE_DATA_TYPE_STRING;
5579         se_val.value.sv_string = (char *)ddi_driver_name(dip);
5580         if (sysevent_add_attr(&ev_attr_list,
5581             DEVFS_DRIVER_NAME, &se_val, SE_SLEEP) != 0) {
5582             sysevent_free_attr(ev_attr_list);
5583             goto fail;
5584         }

5586         se_val.value_type = SE_DATA_TYPE_INT32;
5587         se_val.value.sv_int32 = ddi_get_instance(dip);
5588         if (sysevent_add_attr(&ev_attr_list,
5589             DEVFS_INSTANCE, &se_val, SE_SLEEP) != 0) {
5590             sysevent_free_attr(ev_attr_list);
5591             goto fail;
5592         }

5594     }

5596     if (sysevent_attach_attributes(ev, ev_attr_list) != 0) {
5597         sysevent_free_attr(ev_attr_list);
5598     } else {
5599         (void) log_sysevent(ev, SE_SLEEP, &eid);
5600     }
5601 fail:
5602     sysevent_free(ev);
5603     return (DDI_SUCCESS);
5604 }
_____unchanged_portion_omitted_____

7753 /*
7754  * Allocate and initialize a device id.
7755  */
7756 int
7757 ddi_devid_init(
7758     dev_info_t     *dip,
7759     ushort_t      devid_type,
7760     ushort_t      nbytes,
7761     void           *id,
7762     ddi_devid_t   *ret_devid)
7763 {
7764     impl_devid_t   *i_devid;
7765     int            sz = sizeof (*i_devid) + nbytes - sizeof (char);
7766     int            driver_len;
7767     const char     *driver_name;

7769     switch (devid_type) {
7770     case DEVID_SCSI3_WWN:
7771         /*FALLTHRU*/
7772     case DEVID_SCSI_SERIAL:
7773         /*FALLTHRU*/
7774     case DEVID_ATA_SERIAL:
7775         /*FALLTHRU*/

```

```

7776     case DEVID_ENCAP:
7777         if (nbytes == 0)
7778             return (DDI_FAILURE);
7779         if (id == NULL)
7780             return (DDI_FAILURE);
7781         break;
7782     case DEVID_FAB:
7783         if (nbytes != 0)
7784             return (DDI_FAILURE);
7785         if (id != NULL)
7786             return (DDI_FAILURE);
7787         nbytes = sizeof (int) +
7788             sizeof (struct timeval32) + sizeof (short);
7789         sz += nbytes;
7790         break;
7791     default:
7792         return (DDI_FAILURE);
7793     }

7795     i_devid = kmem_zalloc(sz, KM_SLEEP);
7799     if ((i_devid = kmem_zalloc(sz, KM_SLEEP)) == NULL)
7800         return (DDI_FAILURE);

7797     i_devid->did_magic_hi = DEVID_MAGIC_MSB;
7798     i_devid->did_magic_lo = DEVID_MAGIC_LSB;
7799     i_devid->did_rev_hi = DEVID_REV_MSB;
7800     i_devid->did_rev_lo = DEVID_REV_LSB;
7801     DEVID_FORMTYPE(i_devid, devid_type);
7802     DEVID_FORMLEN(i_devid, nbytes);

7804     /* Fill in driver name hint */
7805     driver_name = ddi_driver_name(dip);
7806     driver_len = strlen(driver_name);
7807     if (driver_len > DEVID_HINT_SIZE) {
7808         /* Pick up last four characters of driver name */
7809         driver_name += driver_len - DEVID_HINT_SIZE;
7810         driver_len = DEVID_HINT_SIZE;
7811     }

7813     bcopy(driver_name, i_devid->did_driver, driver_len);

7815     /* Fill in id field */
7816     if (devid_type == DEVID_FAB) {
7817         char            *cp;
7818         uint32_t        hostid;
7819         struct timeval32 timestamp32;
7820         int             i;
7821         int             *ip;
7822         short           gen;

7824         /* increase the generation number */
7825         mutex_enter(&devid_gen_mutex);
7826         gen = devid_gen_number++;
7827         mutex_exit(&devid_gen_mutex);

7829         cp = i_devid->did_id;

7831         /* Fill in host id (big-endian byte ordering) */
7832         hostid = zone_get_hostid(NULL);
7833         *cp++ = hibyte(hostid);
7834         *cp++ = lobyte(hostid);
7835         *cp++ = hibyte(loword(hostid));
7836         *cp++ = lobyte(loword(hostid));

7838         /*
7839         * Fill in timestamp (big-endian byte ordering)

```

```

7840         *
7841         * (Note that the format may have to be changed
7842         * before 2038 comes around, though it's arguably
7843         * unique enough as it is..)
7844         */
7845         uniqttime32(&timestamp32);
7846         ip = (int *)&timestamp32;
7847         for (i = 0;
7848             i < sizeof (timestamp32) / sizeof (int); i++, ip++) {
7849             int         val;
7850             val = *ip;
7851             *cp++ = hibyte(hiword(val));
7852             *cp++ = lobyte(hiword(val));
7853             *cp++ = hibyte(loword(val));
7854             *cp++ = lobyte(loword(val));
7855         }

7857         /* fill in the generation number */
7858         *cp++ = hibyte(gen);
7859         *cp++ = lobyte(gen);
7860     } else
7861         bcopy(id, i_devid->did_id, nbytes);

7863     /* return device id */
7864     *ret_devid = (ddi_devid_t)i_devid;
7865     return (DDI_SUCCESS);
7866 }
_____unchanged_portion_omitted_

```


new/usr/src/uts/common/rpc/sec_gss/svc_rpcsec_gss.c

1

49471 Thu Oct 23 11:05:00 2014

new/usr/src/uts/common/rpc/sec_gss/svc_rpcsec_gss.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

unchanged_portion_omitted

```
1840 /*
1841  * Add retransmit entry to the context cache entry for a new xid.
1842  * If there is already an entry, delete it before adding the new one.
1843  */
1844 static void retrans_add(client, xid, result)
1845     svc_rpc_gss_data *client;
1846     uint32_t xid;
1847     rpc_gss_init_res *result;
1848 {
1849     retrans_entry *rdata;
1851     if (client->retrans_data && client->retrans_data->xid == xid)
1852         return;
1854     rdata = kmem_zalloc(sizeof (*rdata), KM_SLEEP);
1856     if (rdata == NULL)
1857         return;
1859     rdata->xid = xid;
1860     rdata->result = *result;
1861     if (result->token.length != 0) {
1862         GSS_DUP_BUFFER(rdata->result.token, result->token);
1863     }
1864     if (client->retrans_data)
1865         retrans_del(client);
1866     client->retrans_data = rdata;
1867 }
unchanged_portion_omitted
```

new/usr/src/uts/i86pc/io/gfx_private/gfxp_devmap.c

1

```
*****
4481 Thu Oct 23 11:05:00 2014
new/usr/src/uts/i86pc/io/gfx_private/gfxp_devmap.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */

22 /*
23  * Copyright 2007 Sun Microsystems, Inc. All rights reserved.
24  * Use is subject to license terms.
25 */

27 #pragma ident      "%Z%M% %I%      %E% SMI"

27 #include <sys/debug.h>
28 #include <sys/types.h>
29 #include <sys/param.h>
30 #include <sys/time.h>
31 #include <sys/buf.h>
32 #include <sys/errno.h>
33 #include <sys/system.h>
34 #include <sys/conf.h>
35 #include <sys/signal.h>
36 #include <vm/page.h>
37 #include <vm/as.h>
38 #include <vm/hat.h>
39 #include <vm/seg.h>
40 #include <vm/seg_dev.h>
41 #include <vm/hat_i86.h>
42 #include <sys/ddi.h>
43 #include <sys/devops.h>
44 #include <sys/sunddi.h>
45 #include <sys/ddi_impldefs.h>
46 #include <sys/fs/snode.h>
47 #include <sys/pci.h>
48 #include <sys/vmsystem.h>
49 #include <sys/int_fmtio.h>
50 #include "gfx_private.h"

52 #ifdef __xpv
53 #include <sys/hypervisor.h>
54 #endif

56 /*
57  * Create a dummy ddi_umem_cookie given to gfxp_devmap_umem_setup().
58  */
```

new/usr/src/uts/i86pc/io/gfx_private/gfxp_devmap.c

2

```
59 ddi_umem_cookie_t
60 gfxp_umem_cookie_init(caddr_t kva, size_t size)
61 {
62     struct ddi_umem_cookie *umem_cookie;

64     umem_cookie = kmem_zalloc(sizeof (struct ddi_umem_cookie), KM_SLEEP);

68     if (umem_cookie == NULL)
69         return (NULL);

66     umem_cookie->cvaddr = kva;
67     umem_cookie->type = KMEM_NON_PAGEABLE;
68     umem_cookie->size = size;

70     return ((ddi_umem_cookie_t *)umem_cookie);
71 }
_____unchanged_portion_omitted_____
```

```

*****
8026 Thu Oct 23 11:05:00 2014
new/usr/src/uts/i86pc/io/gfx_private/gfxp_pci.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License, Version 1.0 only
6  * (the "License").  You may not use this file except in compliance
7  * with the License.
8  *
9  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
10 * or http://www.opensolaris.org/os/licensing.
11 * See the License for the specific language governing permissions
12 * and limitations under the License.
13 *
14 * When distributing Covered Code, include this CDDL HEADER in each
15 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
16 * If applicable, add the following below this CDDL HEADER, with the
17 * fields enclosed by brackets "[]" replaced with your own identifying
18 * information: Portions Copyright [yyyy] [name of copyright owner]
19 *
20 * CDDL HEADER END
21 */
22 /*
23 * Copyright 2005 Sun Microsystems, Inc.  All rights reserved.
24 * Use is subject to license terms.
25 */

27 #pragma ident      "%Z%M% %I%      %E% SMI"

27 #include <sys/debug.h>
28 #include <sys/types.h>
29 #include <sys/param.h>
30 #include <sys/time.h>
31 #include <sys/buf.h>
32 #include <sys/errno.h>
33 #include <sys/system.h>
34 #include <sys/conf.h>
35 #include <sys/signal.h>
36 #include <sys/file.h>
37 #include <sys/uio.h>
38 #include <sys/ioctl.h>
39 #include <sys/map.h>
40 #include <sys/proc.h>
41 #include <sys/user.h>
42 #include <sys/mman.h>
43 #include <sys/cred.h>
44 #include <sys/open.h>
45 #include <sys/stat.h>
46 #include <sys/utsname.h>
47 #include <sys/kmem.h>
48 #include <sys/cmn_err.h>
49 #include <sys/vnode.h>
50 #include <vm/page.h>
51 #include <vm/as.h>
52 #include <vm/hat.h>
53 #include <vm/seg.h>
54 #include <sys/ddi.h>
55 #include <sys/devops.h>
56 #include <sys/sunddi.h>
57 #include <sys/ddi_impldefs.h>
58 #include <sys/fs/snode.h>

```

```

59 #include <sys/pci.h>
60 #include <sys/modctl.h>
61 #include <sys/uio.h>
62 #include <sys/visual_io.h>
63 #include <sys/fbio.h>
64 #include <sys/ddidmareq.h>
65 #include <sys/tnf_probe.h>
66 #include <sys/kstat.h>
67 #include <sys/callb.h>
68 #include <sys/pci_cfgspace.h>
69 #include "gfx_private.h"

71 typedef struct gfxp_pci_bsf {
72     uint16_t     vendor;
73     uint16_t     device;
74     uint8_t      bus;
75     uint8_t      slot;
76     uint8_t      function;
77     uint8_t      found;
78     dev_info_t   *dip;
79 } gfxp_pci_bsf_t;
_____ unchanged_portion_omitted _____

152 gfxp_acc_handle_t
153 gfxp_pci_init_handle(uint8_t bus, uint8_t slot, uint8_t function,
154                     uint16_t *vendor, uint16_t *device)
155 {
156     dev_info_t     *dip;
157     gfxp_pci_bsf_t *pci_bsf;

159     /*
160      * Find a PCI device based on its address, and return a unique handle
161      * to be used in subsequent calls to read from or write to the config
162      * space of this device.
163      */

165     pci_bsf = kmem_zalloc(sizeof (gfxp_pci_bsf_t), KM_SLEEP);
166     if ((pci_bsf = kmem_zalloc(sizeof (gfxp_pci_bsf_t), KM_SLEEP))
167         == NULL) {
168         return (NULL);
169     }

167     pci_bsf->bus = bus;
168     pci_bsf->slot = slot;
169     pci_bsf->function = function;

171     ddi_walk_devs(dden_root_node(), gfxp_pci_find_bsf, pci_bsf);

173     if (pci_bsf->found) {
174         dip = pci_bsf->dip;

176         if (vendor) *vendor = pci_bsf->vendor;
177         if (device) *device = pci_bsf->device;
178     } else {
179         dip = NULL;
180         if (vendor) *vendor = 0x0000;
181         if (device) *device = 0x0000;
182     }

184     kmem_free(pci_bsf, sizeof (gfxp_pci_bsf_t));

186     return ((gfxp_acc_handle_t)dip);
187 }
_____ unchanged_portion_omitted _____

318 int

```

```
319 gfxp_pci_device_present(uint16_t vendor, uint16_t device)
320 {
321     gfxp_pci_bsf_t *pci_bsf;
322     int rv;
323
324     /*
325      * Find a PCI device based on its device and vendor id.
326      */
327
328     pci_bsf = kmem_zalloc(sizeof (gfxp_pci_bsf_t), KM_SLEEP);
329     if ((pci_bsf = kmem_zalloc(sizeof (gfxp_pci_bsf_t), KM_SLEEP)) == NULL)
330         return (0);
331
332     pci_bsf->vendor = vendor;
333     pci_bsf->device = device;
334     ddi_walk_devs(ddi_root_node(), gfxp_pci_find_vd, pci_bsf);
335
336     if (pci_bsf->found) {
337         rv = 1;
338     } else {
339         rv = 0;
340     }
341
342     kmem_free(pci_bsf, sizeof (gfxp_pci_bsf_t));
343     return (rv);
344 }
345
346 _____unchanged_portion_omitted_____
```

```

*****
44789 Thu Oct 23 11:05:01 2014
new/usr/src/uts/intel/io/dktp/dcdev/dadk.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

716 /* ARGSUSED */
717 int
718 dadk_ioctl(opaque_t objp, dev_t dev, int cmd, intptr_t arg, int flag,
719          cred_t *cred_p, int *rval_p)
720 {
721     struct dadk *dadkp = (struct dadk *)objp;

722
723     switch (cmd) {
724     case DKIOCGETDEF:
725         {
726             struct buf *bp;
727             int err, head;
728             unsigned char *secbuf;
729             STRUCT_DECL(defect_header, adh);

730
731             STRUCT_INIT(adh, flag & FMODELS);

732
733             /*
734              * copyin header ...
735              * yields head number and buffer address
736              */
737             if (ddi_copyin((caddr_t)arg, STRUCT_BUF(adh), STRUCT_SIZE(adh),
738                          flag))
739                 return (EFAULT);
740             head = STRUCT_FGET(adh, head);
741             if (head < 0 || head >= dadkp->dad_phyg.g_head)
742                 return (ENXIO);
743             secbuf = kmem_zalloc(NBPSCTR, KM_SLEEP);

744             if (!secbuf)
745                 return (ENOMEM);
746             bp = getrbuf(KM_SLEEP);
747             if (!bp) {
748                 kmem_free(secbuf, NBPSCTR);
749                 return (ENOMEM);
750             }

751
752             bp->b_edev = dev;
753             bp->b_dev = cmpdev(dev);
754             bp->b_flags = B_BUSY;
755             bp->b_resid = 0;
756             bp->b_bcount = NBPSCTR;
757             bp->b_un.b_addr = (caddr_t)secbuf;
758             bp->b_blkno = head; /* I had to put it somewhere! */
759             bp->b_forw = (struct buf *)dadkp;
760             bp->b_back = (struct buf *)DCMD_GETDEF;

761
762             mutex_enter(&dadkp->dad_cmd_mutex);
763             dadkp->dad_cmd_count++;
764             mutex_exit(&dadkp->dad_cmd_mutex);
765             FLC_ENQUEUE(dadkp->dad_flcobj, bp);
766             err = biowait(bp);
767             if (!err) {
768                 if (ddi_copyout((caddr_t)secbuf,
769                               STRUCT_FGETP(adh, buffer), NBPSCTR, flag))
770                     err = ENXIO;
771             }
772             kmem_free(secbuf, NBPSCTR);

```

```

768         freerbuf(bp);
769         return (err);
770     }
771     case DIOCTL_RWCMD:
772     {
773         struct dadkio_rwcmd *rwcmdp;
774         int status, rw;

775
776         /*
777          * copied in by cmdk and, if necessary, converted to the
778          * correct datamodel
779          */
780         rwcmdp = (struct dadkio_rwcmd *) (intptr_t) arg;

781
782         /*
783          * handle the complex cases here; we pass these
784          * through to the driver, which will queue them and
785          * handle the requests asynchronously. The simpler
786          * cases, which can return immediately, fail here, and
787          * the request reverts to the dadk_ioctl routine, while
788          * will reroute them directly to the ata driver.
789          */
790         switch (rwcmdp->cmd) {
791         case DADKIO_RWCMD_READ :
792             /* FALLTHROUGH */
793         case DADKIO_RWCMD_WRITE:
794             rw = ((rwcmdp->cmd == DADKIO_RWCMD_WRITE) ?
795                 B_WRITE : B_READ);
796             status = dadk_dk_buf_setup(dadkp,
797                                       (opaque_t)rwcmdp, dev, ((flag & FKIOCTL) ?
798                                         UIO_SYSSPACE : UIO_USERSPACE), rw);
799             return (status);
800         default:
801             return (EINVAL);
802         }
803     }
804     case DKIOC_UPDATEFW:

805         /*
806          * Require PRIV_ALL privilege to invoke DKIOC_UPDATEFW
807          * to protect the firmware update from malicious use
808          */
809         if (PRIV_POLICY(cred_p, PRIV_ALL, B_FALSE, EPERM, NULL) != 0)
810             return (EPERM);
811         else
812             return (dadk_ctl_ioctl(dadkp, cmd, arg, flag));

813
814     case DKIOCFLUSHWRITECACHE:
815     {
816         struct buf *bp;
817         int err = 0;
818         struct dk_callback *dkc = (struct dk_callback *) arg;
819         struct cmpkt *pktp;
820         int is_sync = 1;

821
822         mutex_enter(&dadkp->dad_mutex);
823         if (dadkp->dad_noflush || !dadkp->dad_wce) {
824             err = dadkp->dad_noflush ? ENOTSUP : 0;
825             mutex_exit(&dadkp->dad_mutex);
826             /*
827              * If a callback was requested: a
828              * callback will always be done if the
829              * caller saw the DKIOCFLUSHWRITECACHE
830              * ioctl return 0, and never done if the
831              * caller saw the ioctl return an error.
832              */
833         }

```

```

834         if ((flag & FKIOCTL) && dkc != NULL &&
835             dkc->dkc_callback != NULL) {
836             (*dkc->dkc_callback)(dkc->dkc_cookie,
837                 err);
838             /*
839              * Did callback and reported error.
840              * Since we did a callback, ioctl
841              * should return 0.
842              */
843             err = 0;
844         }
845         return (err);
846     }
847     mutex_exit(&dadkp->dad_mutex);
848
849     bp = getrbuf(KM_SLEEP);
850
851     bp->b_eved = dev;
852     bp->b_dev = cmpdev(dev);
853     bp->b_flags = B_BUSY;
854     bp->b_resid = 0;
855     bp->b_bcount = 0;
856     SET_BP_SEC(bp, 0);
857
858     if ((flag & FKIOCTL) && dkc != NULL &&
859         dkc->dkc_callback != NULL) {
860         struct dk_callback *dkc2 =
861             (struct dk_callback *)kmem_zalloc(
862                 sizeof (struct dk_callback), KM_SLEEP);
863
864         bcopy(dkc, dkc2, sizeof (*dkc2));
865         bp->b_private = dkc2;
866         bp->b_iodone = dadk_flushdone;
867         is_sync = 0;
868     }
869
870     /*
871      * Setup command pkt
872      * dadk_pktprep() can't fail since DDI_DMA_SLEEP set
873      */
874     pktp = dadk_pktprep(dadkp, NULL, bp,
875         dadk_iodone, DDI_DMA_SLEEP, NULL);
876
877     pktp->cp_time = DADK_FLUSH_CACHE_TIME;
878
879     *((char *) (pktp->cp_cdbp)) = DCMD_FLUSH_CACHE;
880     pktp->cp_byteleft = 0;
881     pktp->cp_private = NULL;
882     pktp->cp_seclen = 0;
883     pktp->cp_srtsec = -1;
884     pktp->cp_bytexfer = 0;
885
886     CTL_IOSETUP(dadkp->dad_ctlobjp, pktp);
887
888     mutex_enter(&dadkp->dad_cmd_mutex);
889     dadkp->dad_cmd_count++;
890     mutex_exit(&dadkp->dad_cmd_mutex);
891     FLC_ENQUE(dadkp->dad_flcobjp, bp);
892
893     if (is_sync) {
894         err = biowait(bp);
895         freerbuf(bp);
896     }
897     return (err);
898 }
899 default:

```

```

900         if (!dadkp->dad_rmb)
901             return (dadk_ctl_ioctl(dadkp, cmd, arg, flag));
902     }
903
904     switch (cmd) {
905     case CDROMSTOP:
906         return (dadk_rmb_ioctl(dadkp, DCMD_STOP_MOTOR, 0,
907             0, DADK_SILENT));
908     case CDROMSTART:
909         return (dadk_rmb_ioctl(dadkp, DCMD_START_MOTOR, 0,
910             0, DADK_SILENT));
911     case DKIOCLLOCK:
912         return (dadk_rmb_ioctl(dadkp, DCMD_LOCK, 0, 0, DADK_SILENT));
913     case DKIOCNLOCK:
914         return (dadk_rmb_ioctl(dadkp, DCMD_UNLOCK, 0, 0, DADK_SILENT));
915     case DKIOEJECT:
916     case CDROMEJECT:
917         {
918             int ret;
919
920             if (ret = dadk_rmb_ioctl(dadkp, DCMD_UNLOCK, 0, 0,
921                 DADK_SILENT)) {
922                 return (ret);
923             }
924             if (ret = dadk_rmb_ioctl(dadkp, DCMD_EJECT, 0, 0,
925                 DADK_SILENT)) {
926                 return (ret);
927             }
928             mutex_enter(&dadkp->dad_mutex);
929             dadkp->dad_iostate = DKIO_EJECTED;
930             cv_broadcast(&dadkp->dad_state_cv);
931             mutex_exit(&dadkp->dad_mutex);
932
933             return (0);
934         }
935     default:
936         return (ENOTTY);
937     }
938     /*
939      * cdrom audio commands
940      */
941     case CDROMPAUSE:
942         cmd = DCMD_PAUSE;
943         break;
944     case CDROMRESUME:
945         cmd = DCMD_RESUME;
946         break;
947     case CDROMPLAYMSF:
948         cmd = DCMD_PLAYMSF;
949         break;
950     case CDROMPLAYTRKIND:
951         cmd = DCMD_PLAYTRKIND;
952         break;
953     case CDROMREADTOCHDR:
954         cmd = DCMD_READTOCHDR;
955         break;
956     case CDROMREADTOCENTRY:
957         cmd = DCMD_READTOCENT;
958         break;
959     case CDROMVOLCTRL:
960         cmd = DCMD_VOLCTRL;
961         break;
962     case CDROMSUBCHNL:
963         cmd = DCMD_SUBCHNL;
964         break;
965     case CDROMREADMODE2:

```

```
966         cmd = DCMD_READMODE2;
967         break;
968     case CDROMREADMODEL:
969         cmd = DCMD_READMODE1;
970         break;
971     case CDROMREADOFFSET:
972         cmd = DCMD_READOFFSET;
973         break;
974     }
975     return (dadk_rmb_ioctl(dadkp, cmd, arg, flag, 0));
976 }
```

unchanged portion omitted

```
1681 static int
1682 dadk_rmb_ioctl(struct dadk *dadkp, int cmd, intptr_t arg, int flags, int silent)
```

```
1684 {
1685     struct buf *bp;
1686     int err;
1687     struct cmpkt *pkt;
1688
1689     bp = getrbuf(KM_SLEEP);
1690     if ((bp = getrbuf(KM_SLEEP)) == NULL) {
1691         return (ENOMEM);
1692     }
1693     pkt = dadk_pktprep(dadkp, NULL, bp, dadk_rmb_iodone, NULL, NULL);
1694     if (!pkt) {
1695         freerbuf(bp);
1696         return (ENOMEM);
1697     }
1698     bp->b_back = (struct buf *)arg;
1699     bp->b_forw = (struct buf *)dadkp->dad_flcobjp;
1700     pkt->cp_passthru = (opaque_t)(intptr_t)silent;
1701
1702     err = dadk_ctl_ioctl(dadkp, cmd, (uintptr_t)pkt, flags);
1703     freerbuf(bp);
1704     gda_free(dadkp->dad_ctlobjp, pkt, NULL);
1705     return (err);
1706 }
```

unchanged portion omitted

new/usr/src/uts/intel/io/drm/i915_gem.c

1

79987 Thu Oct 23 11:05:01 2014

new/usr/src/uts/intel/io/drm/i915_gem.c

5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP

5254 getrbuf won't fail with KM_SLEEP

unchanged_portion_omitted

```
2444 static int
2445 i915_gem_object_get_page_list(struct drm_gem_object *obj)
2446 {
2447     struct drm_i915_gem_object *obj_priv = obj->driver_private;
2448     caddr_t va;
2449     long i;

2451     if (obj_priv->page_list)
2452         return 0;
2453     pgcnt_t np = btop(obj->size);

2455     obj_priv->page_list = kmem_zalloc(np * sizeof(caddr_t), KM_SLEEP);
2456     if (obj_priv->page_list == NULL) {
2457         DRM_ERROR("Failed to allocate page list\n");
2458         return ENOMEM;
2459     }

2457     for (i = 0, va = obj->kaddr; i < np; i++, va += PAGE_SIZE) {
2458         obj_priv->page_list[i] = va;
2459     }
2460     return 0;
2461 }
```

unchanged_portion_omitted


```

*****
32818 Thu Oct 23 11:05:01 2014
new/usr/src/uts/intel/io/heci/heci_init.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
unchanged_portion_omitted_

758 /*
759 * allocate_me_clients_storage - allocate storage for me clients
760 *
761 * @dev: Device object for our driver
762 *
763 * @return 0 on success, <0 on failure.
764 */
765 static int
766 allocate_me_clients_storage(struct iamt_heci_device *dev)
767 {
768     struct heci_me_client *clients;
769     struct heci_me_client *client;
770     uint8_t num, i, j;
771     int err;

773     if (dev->num_heci_me_clients == 0)
774         return (0);

776     mutex_enter(&dev->device_lock);
777     if (dev->me_clients) {
778         kmem_free(dev->me_clients, dev->num_heci_me_clients*
779             sizeof (struct heci_me_client));
780         dev->me_clients = NULL;
781     }
782     mutex_exit(&dev->device_lock);

784     /* allocate storage for ME clients representation */
785     clients = kmem_zalloc(dev->num_heci_me_clients*
786         sizeof (struct heci_me_client), KM_SLEEP);
787     if (!clients) {
788         DBG("memory allocation for ME clients failed.\n");
789         return (-ENOMEM);
790     }

788     mutex_enter(&dev->device_lock);
789     dev->me_clients = clients;
790     mutex_exit(&dev->device_lock);

792     num = 0;
793     for (i = 0; i < sizeof (dev->heci_me_clients); i++) {
794         for (j = 0; j < 8; j++) {
795             if ((dev->heci_me_clients[i] & (1 << j)) != 0) {
796                 client = &dev->me_clients[num];
797                 client->client_id = (i * 8) + j;
798                 client->flow_ctrl_creds = 0;
799                 err = host_client_properties(dev, client);
800                 if (err != 0) {
801                     mutex_enter(&dev->device_lock);
802                     kmem_free(dev->me_clients,
803                         dev->num_heci_me_clients*
804                         sizeof (struct heci_me_client));
805                     dev->me_clients = NULL;
806                     mutex_exit(&dev->device_lock);
807                     return (err);
808                 }
809                 num++;
810             }
811         }
}

```

```

812     }
814     return (0);
815 }
unchanged_portion_omitted_

1038 /*
1039 * heci_alloc_file_private - allocates a private file structure and set it up.
1040 * @file: the file structure
1041 *
1042 * @return The allocated file or NULL on failure
1043 */
1044 struct heci_file_private *
1045 heci_alloc_file_private(struct heci_file *file)
1046 {
1047     struct heci_file_private *priv;

1049     priv = kmem_zalloc(sizeof (struct heci_file_private), KM_SLEEP);
1054     if (!priv)
1055         return (NULL);

1051     heci_init_file_private(priv, file);

1053     return (priv);
1054 }
unchanged_portion_omitted_

1071 /*
1072 * heci_disconnect_host_client - send disconnect message to fw from host
1073 * client.
1074 *
1075 * @dev: Device object for our driver
1076 * @file_ext: private data of the file object
1077 *
1078 * @return 0 on success, <0 on failure.
1079 */
1080 int
1081 heci_disconnect_host_client(struct iamt_heci_device *dev,
1082     struct heci_file_private *file_ext)
1083 {
1084     int rets, err;
1085     long timeout = 15; /* 15 seconds */
1086     struct heci_cb_private *priv_cb;
1087     clock_t delta = (clock_t)(timeout * HZ);

1089     if (!dev || !file_ext)
1090         return (-ENODEV);

1092     if (file_ext->state != HECI_FILE_DISCONNECTING)
1093         return (0);

1095     priv_cb = kmem_zalloc(sizeof (struct heci_cb_private), KM_SLEEP);
1102     if (!priv_cb)
1103         return (-ENOMEM);

1097     LIST_INIT_HEAD(&priv_cb->cb_list);
1098     priv_cb->file_private = file_ext;
1099     priv_cb->major_file_operations = HECI_CLOSE;
1100     mutex_enter(&dev->device_lock);
1101     if (dev->host_buffer_is_empty) {
1102         dev->host_buffer_is_empty = 0;
1103         if (heci_disconnect(dev, file_ext)) {
1104             list_add_tail(&priv_cb->cb_list,
1105                 &dev->ctrl_rd_list.heci_cb.cb_list);
1106         } else {
1107             mutex_exit(&dev->device_lock);
}

```

```
1108             rets = -ENODEV;
1109             DBG("failed to call heci_disconnect.\n");
1110             goto free;
1111         }
1112     } else {
1113         DBG("add disconnect cb to control write list\n");
1114         list_add_tail(&priv_cb->cb_list,
1115                     &dev->ctrl_wr_list.heci_cb.cb_list);
1116     }
1117
1118     err = 0;
1119     while (err != -1 &&
1120           (HECI_FILE_DISCONNECTED != file_ext->state)) {
1121
1122         err = cv_reltimedwait(&dev->wait_recvd_msg, &dev->device_lock,
1123                             delta, TR_CLOCK_TICK);
1124     }
1125     mutex_exit(&dev->device_lock);
1126
1127     if (HECI_FILE_DISCONNECTED == file_ext->state) {
1128         rets = 0;
1129         DBG("successfully disconnected from fw client."
1130            " me_client_id:%d, host_client_id:%d\n",
1131            file_ext->me_client_id,
1132            file_ext->host_client_id);
1133     } else {
1134         rets = -ENODEV;
1135         if (HECI_FILE_DISCONNECTED != file_ext->state)
1136             DBG("wrong status client disconnect.\n");
1137
1138         if (err)
1139             DBG("wait failed disconnect err=%08x\n", err);
1140
1141         DBG("failed to disconnect from fw client.\n"
1142            " me_client_id:%d, host_client_id:%d\n",
1143            file_ext->me_client_id,
1144            file_ext->host_client_id);
1145     }
1146
1147     mutex_enter(&dev->device_lock);
1148     heci_flush_list(&dev->ctrl_rd_list, file_ext);
1149     heci_flush_list(&dev->ctrl_wr_list, file_ext);
1150     mutex_exit(&dev->device_lock);
1151 free:
1152     heci_free_cb_private(priv_cb);
1153     return (rets);
1154 }
unchanged_portion_omitted
```

```
*****
```

```
25971 Thu Oct 23 11:05:01 2014
```

```
new/usr/src/uts/intel/os/fmsmb.c
```

```
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
```

```
5254 getrbuf won't fail with KM_SLEEP
```

```
*****
```

```
unchanged_portion_omitted
```

```
631 /*
632  * go throught the type 2 structure contained_ids looking for
633  * the type 4 which has strand_apicid == this strand_apicid
634  */
635 static int
636 find_matching_proc(smbios_hdl_t *shp, uint_t strand_apicid,
637                  uint16_t bb_id, uint16_t proc_hdl, int is_proc)
638 {
639     int n;
640     const smb_struct_t *sp;
641     smbios_bboard_t bb;
642     uint_t cont_count, cont_len;
643     uint16_t cont_id;
644     id_t *cont_hdl = NULL;
645     int rc;

648     (void) smbios_info_bboard(shp, bb_id, &bb);
649     cont_count = (uint_t)bb.smbb_contn;
650     if (cont_count == 0)
651         return (0);

653     cont_len = sizeof (id_t);
654     cont_hdl = kmem_zalloc(cont_count * cont_len, KM_SLEEP);
655     if (cont_hdl == NULL)
656         return (0);

656     rc = smbios_info_contains(shp, bb_id, cont_count, cont_hdl);
657     if (rc > SMB_CONT_MAX) {
658         kmem_free(cont_hdl, cont_count * cont_len);
659         return (0);
660     }
661     cont_count = MIN(rc, cont_count);

663     for (n = 0; n < cont_count; n++) {
664         cont_id = (uint16_t)cont_hdl[n];
665         sp = smb_lookup_id(shp, cont_id);
666         if (sp->smbst_hdr->smbh_type == SMB_TYPE_PROCESSOR) {
667             if (is_proc) {
668                 if (find_matching_apic(shp, cont_id,
669                                     strand_apicid) {
670                     kmem_free(cont_hdl,
671                               cont_count * cont_len);
672                     return (1);
673                 }
674             } else {
675                 if (cont_id == proc_hdl) {
676                     kmem_free(cont_hdl,
677                               cont_count * cont_len);
678                     return (1);
679                 }
680             }
681         }
682     }
683     if (cont_hdl != NULL)
684         kmem_free(cont_hdl, cont_count * cont_len);

686     return (0);
```

```
687 }
```

```
unchanged_portion_omitted
```

new/usr/src/uts/sun4/io/pcicfg.c

1

```
*****
182715 Thu Oct 23 11:05:01 2014
new/usr/src/uts/sun4/io/pcicfg.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
_____unchanged_portion_omitted_____

5783 #ifndef PCICFG_INTERPRET_FCODE
5784 static int
5785 pcicfg_load_fcode(dev_info_t *dip, uint_t bus, uint_t device, uint_t func,
5786     uint16_t vendor_id, uint16_t device_id, uchar_t **fcode_addr,
5787     int *fcode_size, int rom_paddr, int rom_size)
5788 {
5789     pci_regspec_t      p;
5790     int                pci_data;
5791     int                start_of_fcode;
5792     int                image_length;
5793     int                code_type;
5794     ddi_acc_handle_t   h;
5795     ddi_device_acc_attr_t acc;
5796     uint8_t            *addr;
5797     int8_t             image_not_found, indicator;
5798     uint16_t           vendor_id_img, device_id_img;
5799     int16_t            rom_sig;
5800 #ifdef DEBUG
5801     int i;
5802 #endif

5804     DEBUG4("pcicfg_load_fcode() - "
5805         "bus %x device %x func=%x rom_paddr=%lx\n",
5806         bus, device, func, rom_paddr);
5807     DEBUG2("pcicfg_load_fcode() - vendor_id=%x device_id=%x\n",
5808         vendor_id, device_id);

5810     *fcode_size = 0;
5811     *fcode_addr = NULL;

5813     acc.devacc_attr_version = DDI_DEVICE_ATTR_V0;
5814     acc.devacc_attr_endian_flags = DDI_STRUCTURE_LE_ACC;
5815     acc.devacc_attr_dataorder = DDI_STRICTORDER_ACC;

5817     p.pci_phys_hi = PCI_ADDR_MEM32 | PCICFG_MAKE_REG_HIGH(bus, device,
5818         func, PCI_CONF_ROM);

5820     p.pci_phys_mid = 0;
5821     p.pci_phys_low = 0;

5823     p.pci_size_low = rom_size;
5824     p.pci_size_hi = 0;

5826     if (pcicfg_map_phys(dip, &p, (caddr_t *)&addr, &acc, &h) {
5827         DEBUG1("Can Not map in ROM %x\n", p.pci_phys_low);
5828         return (PCICFG_FAILURE);
5829     }

5831     /*
5832     * Walk the ROM to find the proper image for this device.
5833     */
5834     image_not_found = 1;
5835     while (image_not_found) {
5836         DEBUG1("Expansion ROM maps to %lx\n", addr);

5838 #ifdef DEBUG
5839     if (pcicfg_dump_fcode) {
5840         for (i = 0; i < 100; i++)
```

new/usr/src/uts/sun4/io/pcicfg.c

2

```
5841         DEBUG2("ROM 0x%x --> 0x%x\n", i,
5842             ddi_get8(h, (uint8_t *) (addr + i)));
5843     }
5844 #endif

5846     /*
5847     * Some device say they have an Expansion ROM, but do not, so
5848     * for non-21554 devices use peek so we don't panic due to
5849     * accessing non existent memory.
5850     */
5851     if (pcicfg_indirect_map(dip) == DDI_SUCCESS) {
5852         rom_sig = ddi_get16(h,
5853             (uint16_t *) (addr + PCI_ROM_SIGNATURE));
5854     } else {
5855         if (ddi_peek16(dip,
5856             (int16_t *) (addr + PCI_ROM_SIGNATURE), &rom_sig) {
5857             cmn_err(CE_WARN,
5858                 "PCI Expansion ROM is not accessible");
5859             pcicfg_unmap_phys(&h, &p);
5860             return (PCICFG_FAILURE);
5861         }
5862     }

5864     /*
5865     * Validate the ROM Signature.
5866     */
5867     if ((uint16_t)rom_sig != 0xaa55) {
5868         DEBUG1("Invalid ROM Signature %x\n", (uint16_t)rom_sig);
5869         pcicfg_unmap_phys(&h, &p);
5870         return (PCICFG_FAILURE);
5871     }

5873     DEBUG0("Valid ROM Signature Found\n");

5875     start_of_fcode = ddi_get16(h, (uint16_t *) (addr + 2));

5877     pci_data = ddi_get16(h,
5878         (uint16_t *) (addr + PCI_ROM_PCI_DATA_STRUCT_PTR));

5880     DEBUG2("Pointer To PCI Data Structure %x %x\n", pci_data,
5881         addr);

5883     /*
5884     * Validate the PCI Data Structure Signature.
5885     * 0x52494350 = "PCIR"
5886     */

5888     if (ddi_get8(h, (uint8_t *) (addr + pci_data)) != 0x50) {
5889         DEBUG0("Invalid PCI Data Structure Signature\n");
5890         pcicfg_unmap_phys(&h, &p);
5891         return (PCICFG_FAILURE);
5892     }

5894     if (ddi_get8(h, (uint8_t *) (addr + pci_data + 1)) != 0x43) {
5895         DEBUG0("Invalid PCI Data Structure Signature\n");
5896         pcicfg_unmap_phys(&h, &p);
5897         return (PCICFG_FAILURE);
5898     }
5899     if (ddi_get8(h, (uint8_t *) (addr + pci_data + 2)) != 0x49) {
5900         DEBUG0("Invalid PCI Data Structure Signature\n");
5901         pcicfg_unmap_phys(&h, &p);
5902         return (PCICFG_FAILURE);
5903     }
5904     if (ddi_get8(h, (uint8_t *) (addr + pci_data + 3)) != 0x52) {
5905         DEBUG0("Invalid PCI Data Structure Signature\n");
5906         pcicfg_unmap_phys(&h, &p);
```

```

5907         return (PCICFG_FAILURE);
5908     }
5910     /*
5911     * Is this image for this device?
5912     */
5913     vendor_id_img = ddi_get16(h,
5914         (uint16_t *) (addr + pci_data + PCI_PDS_VENDOR_ID));
5915     device_id_img = ddi_get16(h,
5916         (uint16_t *) (addr + pci_data + PCI_PDS_DEVICE_ID));
5918     DEBUG2("This image is for vendor_id=%x device_id=%x\n",
5919         vendor_id_img, device_id_img);
5921     code_type = ddi_get8(h, addr + pci_data + PCI_PDS_CODE_TYPE);
5923     switch (code_type) {
5924     case PCI_PDS_CODE_TYPE_PCAT:
5925         DEBUG0("ROM is of x86/PC-AT Type\n");
5926         break;
5927     case PCI_PDS_CODE_TYPE_OPEN_FW:
5928         DEBUG0("ROM is of Open Firmware Type\n");
5929         break;
5930     default:
5931         DEBUG1("ROM is of Unknown Type 0x%x\n", code_type);
5932         break;
5933     }
5935     if ((vendor_id_img != vendor_id) ||
5936         (device_id_img != device_id) ||
5937         (code_type != PCI_PDS_CODE_TYPE_OPEN_FW)) {
5938         DEBUG0("Firmware Image is not for this device..."
5939             "goto next image\n");
5940         /*
5941         * Read indicator byte to see if there is another
5942         * image in the ROM
5943         */
5944         indicator = ddi_get8(h,
5945             (uint8_t *) (addr + pci_data + PCI_PDS_INDICATOR));
5947         if (indicator != 1) {
5948             /*
5949             * There is another image in the ROM.
5950             */
5951             image_length = ddi_get16(h, (uint16_t *) (addr +
5952                 pci_data + PCI_PDS_IMAGE_LENGTH)) * 512;
5954             addr += image_length;
5955         } else {
5956             /*
5957             * There are no more images.
5958             */
5959             DEBUG0("There are no more images in the ROM\n");
5960             pcicfg_unmap_phys(&h, &p);
5962             return (PCICFG_FAILURE);
5963         }
5964     } else {
5965         DEBUG0("Correct image was found\n");
5966         image_not_found = 0; /* Image was found */
5967     }
5968 }
5970 *fcode_size = (ddi_get8(h, addr + start_of_fcode + 4) << 24) |
5971 (ddi_get8(h, addr + start_of_fcode + 5) << 16) |
5972 (ddi_get8(h, addr + start_of_fcode + 6) << 8) |

```

```

5973         (ddi_get8(h, addr + start_of_fcode + 7));
5975     DEBUG1("Fcode Size %x\n", *fcode_size);
5977     /*
5978     * Allocate page aligned buffer space
5979     */
5980     *fcode_addr = kmem_zalloc(ptob(btopr(*fcode_size)), KM_SLEEP);
5982     if (*fcode_addr == NULL) {
5983         DEBUG0("kmem_zalloc returned NULL\n");
5984         pcicfg_unmap_phys(&h, &p);
5985         return (PCICFG_FAILURE);
5986     }
5988     DEBUG1("Fcode Addr %lx\n", *fcode_addr);
5984     ddi_rep_get8(h, *fcode_addr, addr + start_of_fcode, *fcode_size,
5985         DDI_DEV_AUTOINCR);
5987     pcicfg_unmap_phys(&h, &p);
5989     return (PCICFG_SUCCESS);
5990 }

```

unchanged portion omitted

```

*****
9593 Thu Oct 23 11:05:02 2014
new/usr/src/uts/sun4/io/px/px_mmu.c
5253 kmem_alloc/kmem_zalloc won't fail with KM_SLEEP
5254 getrbuf won't fail with KM_SLEEP
*****
1 /*
2  * CDDL HEADER START
3  *
4  * The contents of this file are subject to the terms of the
5  * Common Development and Distribution License (the "License").
6  * You may not use this file except in compliance with the License.
7  *
8  * You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
9  * or http://www.opensolaris.org/os/licensing.
10 * See the License for the specific language governing permissions
11 * and limitations under the License.
12 *
13 * When distributing Covered Code, include this CDDL HEADER in each
14 * file and include the License file at usr/src/OPENSOLARIS.LICENSE.
15 * If applicable, add the following below this CDDL HEADER, with the
16 * fields enclosed by brackets "[]" replaced with your own identifying
17 * information: Portions Copyright [yyyy] [name of copyright owner]
18 *
19 * CDDL HEADER END
20 */
21 /*
22 * Copyright (c) 2005, 2010, Oracle and/or its affiliates. All rights reserved.
23 * Copyright 2012 Garrett D'Amore <garrett@damore.org>. All rights reserved.
24 */

26 /*
27  * PX mmu initialization and configuration
28  */
29 #include <sys/types.h>
30 #include <sys/kmem.h>
31 #include <sys/async.h>
32 #include <sys/sysmacros.h>
33 #include <sys/sunddi.h>
34 #include <sys/ddi_impldefs.h>
35 #include <sys/vmem.h>
36 #include <sys/machsystem.h> /* lddphys() */
37 #include <sys/iommu_tsb.h>
38 #include "px_obj.h"

40 int
41 px_mmu_attach(px_t *px_p)
42 {
43     dev_info_t      *dip = px_p->px_dip;
44     px_mmu_t        *mmu_p;
45     uint32_t        tsb_i = 0;
46     char             map_name[32];
47     px_dvma_range_prop_t *dvma_prop;
48     int             dvma_prop_len;
49     uint32_t        cache_size, tsb_entries;

51     /*
52      * Allocate mmu state structure and link it to the
53      * px state structure.
54      */
55     mmu_p = kmem_zalloc(sizeof (px_mmu_t), KM_SLEEP);
56     if (mmu_p == NULL)
57         return (DDI_FAILURE);

57     px_p->px_mmu_p = mmu_p;
58     mmu_p->mmu_px_p = px_p;

```

```

59     mmu_p->mmu_inst = ddi_get_instance(dip);

61     /*
62      * Check for "virtual-dma" property that specifies
63      * the DVMA range.
64      */
65     if (ddi_getlongprop(DDI_DEV_T_ANY, dip, DDI_PROP_DONTPASS,
66         "virtual-dma", (caddr_t)&dvma_prop, &dvma_prop_len) !=
67         DDI_PROP_SUCCESS) {

69         DBG(DBG_ATTACH, dip, "Getting virtual-dma failed\n");

71         kmem_free(mmu_p, sizeof (px_mmu_t));
72         px_p->px_mmu_p = NULL;

74         return (DDI_FAILURE);
75     }

77     mmu_p->mmu_dvma_base = dvma_prop->dvma_base;
78     mmu_p->mmu_dvma_end = dvma_prop->dvma_base +
79         (dvma_prop->dvma_len - 1);
80     tsb_entries = MMU_BTOP(dvma_prop->dvma_len);

82     kmem_free(dvma_prop, dvma_prop_len);

84     /*
85      * Setup base and bounds for DVMA and bypass mappings.
86      */
87     mmu_p->mmu_dvma_cache_locks =
88         kmem_zalloc(px_dvma_page_cache_entries, KM_SLEEP);

90     mmu_p->dvma_base_pg = MMU_BTOP(mmu_p->mmu_dvma_base);
91     mmu_p->mmu_dvma_reserve = tsb_entries >> 1;
92     mmu_p->dvma_end_pg = MMU_BTOP(mmu_p->mmu_dvma_end);

94     /*
95      * Create a virtual memory map for dvma address space.
96      * Reserve 'size' bytes of low dvma space for fast track cache.
97      */
98     (void) snprintf(map_name, sizeof (map_name), "%s%d_dvma",
99         ddi_driver_name(dip), ddi_get_instance(dip));

101     cache_size = MMU_PTOB(px_dvma_page_cache_entries *
102         px_dvma_page_cache_clustsz);
103     mmu_p->mmu_dvma_fast_end = mmu_p->mmu_dvma_base +
104         cache_size - 1;

106     mmu_p->mmu_dvma_map = vmem_create(map_name,
107         (void *) (mmu_p->mmu_dvma_fast_end + 1),
108         MMU_PTOB(tsb_entries) - cache_size, MMU_PAGE_SIZE,
109         NULL, NULL, NULL, MMU_PAGE_SIZE, VM_SLEEP);

111     mutex_init(&mmu_p->dvma_debug_lock, NULL, MUTEX_DRIVER, NULL);

113     for (tsb_i = 0; tsb_i < tsb_entries; tsb_i++) {
114         r_addr_t ra = 0;
115         io_attributes_t attr;
116         caddr_t va;

118         if (px_lib_iommu_getmap(px_p->px_dip, PCI_TSBID(0, tsb_i),
119             &attr, &ra) != DDI_SUCCESS)
120             continue;

122         va = (caddr_t) (MMU_PTOB(mmu_p->dvma_base_pg + tsb_i));

124         if (va <= (caddr_t) mmu_p->mmu_dvma_fast_end) {

```

```
125         uint32_t cache_i;
126
127         /*
128          * the va is within the *fast* dvma range; therefore,
129          * lock its fast dvma page cache cluster in order to
130          * both preserve the TTE and prevent the use of this
131          * fast dvma page cache cluster by px_dvma_map_fast().
132          * the lock value 0xFF comes from ldstub().
133          */
134         cache_i = tsb_i / px_dvma_page_cache_clustsz;
135         ASSERT(cache_i < px_dvma_page_cache_entries);
136         mmu_p->mmu_dvma_cache_locks[cache_i] = 0xFF;
137     } else {
138         (void) vmem_xalloc(mmu_p->mmu_dvma_map, MMU_PAGE_SIZE,
139             MMU_PAGE_SIZE, 0, 0, va, va + MMU_PAGE_SIZE,
140             VM_NOSLEEP | VM_BESTFIT | VM_PANIC);
141     }
142 }
143
144     return (DDI_SUCCESS);
145 }
146
147 _____unchanged portion omitted_____
```