

```
*****
12185 Thu Feb 20 18:59:04 2014
new/usr/src/lib/libdladm/common/linkprop.c
2553 mac address should be a dladm link property
*****
_____ unchanged_portion_omitted _____
129 static dld_ioc_macprop_t *i_dladm_buf_alloc_by_name(size_t, datalink_id_t,
130             const char *, uint_t, dladm_status_t *);
131 static dld_ioc_macprop_t *i_dladm_buf_alloc_by_id(size_t, datalink_id_t,
132             mac_prop_id_t, uint_t, dladm_status_t *);
133 static dladm_status_t i_dladm_get_public_prop(dladm_handle_t, datalink_id_t,
134             char *, uint_t, uint_t *, void *, size_t);
136 static dladm_status_t i_dladm_set_private_prop(dladm_handle_t, datalink_id_t,
137             const char *, char **, uint_t, uint_t);
138 static dladm_status_t i_dladm_get_priv_prop(dladm_handle_t, datalink_id_t,
139             const char *, char **, uint_t *, dladm_prop_type_t,
140             uint_t);
141 static dladm_status_t i_dladm_macprop(dladm_handle_t, void *, boolean_t);
142 static const char *dladm_perm2str(uint_t, char *);
143 static link_attr_t *dladm_name2prop(const char *);
144 static link_attr_t *dladm_id2prop(mac_prop_id_t);
146 static pd_getf_t get_zone, get_autopush, get_rate_mod, get_rate,
147             get_speed, get_channel, get_powermode, get_radio,
148             get_duplex, get_link_state, get_binary, get_uint32,
149             get_flowctl, get_maxbw, get_cpus, get_priority,
150             get_tagmode, get_range, get_stp, get_bridge_forward,
151             get_bridge_pvid, get_protection, get_rxrings,
152             get_txrings, get_cntavail, get_macaddr,
153             get_txrings, get_cntavail,
154             get_allowedips, get_allowedcdids, get_pool,
155             get_rings_range, get_linkmode_prop;
156 static pd_setf_t set_zone, set_rate, set_powermode, set_radio,
157             set_public_prop, set_resource, set_stp_prop,
158             set_bridge_forward, set_bridge_pvid, set_macaddr;
159             set_bridge_forward, set_bridge_pvid;
160 static pd_checkf_t check_zone, check_autopush, check_rate, check_hoplimit,
161             check_encaplim, check_uint32, check_maxbw, check_cpus,
162             check_stp_prop, check_bridge_pvid, check_allowedips,
163             check_allowedcdids, check_rings, check_macaddr,
164             check_allowedcdids, check_rings,
165             check_pool, check_prop;
166 struct prop_desc {
167     /*
168     * link property name
169     */
170     char *pd_name;
171     /*
172     * default property value, can be set to { "", NULL }
173     */
174     val_desc_t pd_defval;
175     /*
176     * list of optional property values, can be NULL.
177     *
178     * This is set to non-NULL if there is a list of possible property
179     * values. pd_optval would point to the array of possible values.
180     */
181     val_desc_t *pd_optval;
182 }
```

```
185     /*
186     * count of the above optional property values. 0 if pd_optval is NULL.
187     */
188     uint_t pd_noptval;
189     /*
190     * callback to set link property; set to NULL if this property is
191     * read-only and may be called before or after permanent update; see
192     * flags.
193     */
194     pd_setf_t *pd_set;
195     /*
196     * callback to get modifiable link property
197     */
198     pd_getf_t *pd_getmod;
199     /*
200     * callback to get current link property
201     */
202     pd_getf_t *pd_get;
203     /*
204     * callback to validate link property value, set to NULL if pd_optval
205     * is not NULL. In that case, validate the value by comparing it with
206     * the pd_optval. Return a val_desc_t array pointer if the value is
207     * valid.
208     */
209     pd_checkf_t *pd_check;
210     uint_t pd_flags;
211 #define PD_TEMPONLY 0x1 /* property is temporary only */
212 #define PD_CHECK_ALLOC 0x2 /* alloc vd_val as part of pd_check */
213 #define PD_AFTER_PERM 0x4 /* pd_set after db update; no temporary */
214     /*
215     * indicate link classes this property applies to.
216     */
217     datalink_class_t pd_class;
218     /*
219     * indicate link media type this property applies to.
220     */
221     datalink_media_t pd_dmedia;
222 };
223 /*
224  * Supported link properties enumerated in the prop_table[] array are
225  * computed using the callback functions in that array. To compute the
226  * property value, multiple distinct system calls may be needed (e.g.,
227  * for wifi speed, we need to issue system calls to get desired/supported
228  * rates). The link_attr[] table enumerates the interfaces to the kernel,
229  * and the type/size of the data passed in the user-kernel interface.
230 */
231 static link_attr_t link_attr[] = {
232     { MAC_PROP_DUPLEX, sizeof(link_duplex_t), "duplex" },
233     { MAC_PROP_SPEED, sizeof(uint64_t), "speed" },
234     { MAC_PROP_STATUS, sizeof(link_state_t), "state" },
235     { MAC_PROP_AUTONEG, sizeof(uint8_t), "adv_autoneg_cap" },
236     { MAC_PROP_MTU, sizeof(uint32_t), "mtu" },
237 }
```

```

251     { MAC_PROP_FLOWCTRL,      sizeof (link_flowctrl_t), "flowctrl"},  

253     { MAC_PROP_ZONE,         sizeof (dld_ioc_zid_t), "zone"},  

255     { MAC_PROP_AUTO PUSH,    sizeof (struct dlauto push), "autopush"},  

257     { MAC_PROP_ADV_10GFDX_CAP, sizeof (uint8_t),      "adv_10gfdx_cap"},  

259     { MAC_PROP_EN_10GFDX_CAP, sizeof (uint8_t),      "en_10gfdx_cap"},  

261     { MAC_PROP_ADV_1000FDX_CAP, sizeof (uint8_t),     "adv_1000fdx_cap"},  

263     { MAC_PROP_EN_1000FDX_CAP, sizeof (uint8_t),     "en_1000fdx_cap"},  

265     { MAC_PROP_ADV_1000HDX_CAP, sizeof (uint8_t),     "adv_1000hdx_cap"},  

267     { MAC_PROP_EN_1000HDX_CAP, sizeof (uint8_t),     "en_1000hdx_cap"},  

269     { MAC_PROP_ADV_100FDX_CAP, sizeof (uint8_t),     "adv_100fdx_cap"},  

271     { MAC_PROP_EN_100FDX_CAP, sizeof (uint8_t),     "en_100fdx_cap"},  

273     { MAC_PROP_ADV_100HDX_CAP, sizeof (uint8_t),     "adv_100hdx_cap"},  

275     { MAC_PROP_EN_100HDX_CAP, sizeof (uint8_t),     "en_100hdx_cap"},  

277     { MAC_PROP_ADV_10FDX_CAP,  sizeof (uint8_t),     "adv_10fdx_cap"},  

279     { MAC_PROP_EN_10FDX_CAP,  sizeof (uint8_t),     "en_10fdx_cap"},  

281     { MAC_PROP_ADV_10HDX_CAP, sizeof (uint8_t),     "adv_10hdx_cap"},  

283     { MAC_PROP_EN_10HDX_CAP, sizeof (uint8_t),     "en_10hdx_cap"},  

285     { MAC_PROP_WL_ESSID,      sizeof (wl_linkstatus_t), "essid"},  

287     { MAC_PROP_WL_BSSID,      sizeof (wl_bssid_t),     "bssid"},  

289     { MAC_PROP_WL_BSSTYPE,    sizeof (wl_bss_type_t), "bsstype"},  

291     { MAC_PROP_WL_LINKSTATUS, sizeof (wl_linkstatus_t), "wl_linkstatus"},  

293     /* wl_rates_t has variable length */  

294     { MAC_PROP_WL_DESIRED_RATES, sizeof (wl_rates_t), "desired_rates"},  

296     /* wl_rates_t has variable length */  

297     { MAC_PROP_WL_SUPPORTED_RATES, sizeof (wl_rates_t), "supported_rates"},  

299     { MAC_PROP_WL_AUTH_MODE,   sizeof (wl_authmode_t), "authmode"},  

301     { MAC_PROP_WL_ENCRYPTION, sizeof (wl_encryption_t), "encryption"},  

303     { MAC_PROP_WL_RSSI,       sizeof (wl_rssi_t),      "signal"},  

305     { MAC_PROP_WL_PHY_CONFIG, sizeof (wl_phy_conf_t), "phy_conf"},  

307     { MAC_PROP_WL_CAPABILITY, sizeof (wl_capability_t), "capability"},  

309     { MAC_PROP_WL_WPA,        sizeof (wl_wpa_t),       "wpa"},  

311     /* wl_wpa_ess_t has variable length */  

312     { MAC_PROP_WL_SCANRESULTS, sizeof (wl_wpa_ess_t), "scan_results"},  

314     { MAC_PROP_WL_POWER_MODE, sizeof (wl_ps_mode_t), "powermode"},  

316     { MAC_PROP_WL_RADIO,      sizeof (dladm_wlan_radio_t), "wl_radio"},  


```

```

318     { MAC_PROP_WL_ESS_LIST,  sizeof (wl_ess_list_t), "wl_ess_list"},  

320     { MAC_PROP_WL_KEY_TAB,   sizeof (wl_wep_key_tab_t), "wl_wep_key"},  

322     { MAC_PROP_WL_CREATE_IBSS, sizeof (wl_create_ibss_t), "createibss"},  

324     /* wl_wpa_ie_t has variable length */  

325     { MAC_PROP_WL_SETOPTIE,  sizeof (wl_wpa_ie_t), "set_ie"},  

327     { MAC_PROP_WL_DELKEY,    sizeof (wl_del_key_t), "wpa_del_key"},  

329     { MAC_PROP_WL_KEY,      sizeof (wl_key_t),      "wl_key"},  

331     { MAC_PROP_WL_MLME,     sizeof (wl_mlme_t),     "mlme"},  

333     { MAC_PROP_TAGMODE,    sizeof (link_tagmode_t), "tagmode"},  

335     { MAC_PROP_IPTUN_HOPLIMIT, sizeof (uint32_t), "hoplimit"},  

337     { MAC_PROP_IPTUN_ENCAPLIMIT, sizeof (uint32_t), "encaplimit"},  

339     { MAC_PROP_PVID,        sizeof (uint16_t),     "default_tag"},  

341     { MAC_PROP_LLIMIT,      sizeof (uint32_t),     "learn_limit"},  

343     { MAC_PROP_LDECAY,      sizeof (uint32_t),     "learn_decay"},  

345     { MAC_PROP_RESOURCE,   sizeof (mac_resource_props_t), "resource"},  

347     { MAC_PROP_RESOURCE_EFFECTIVE, sizeof (mac_resource_props_t), "resource-effective"},  

348  

350     { MAC_PROP_RXRINGS RANGE, sizeof (mac_propval_range_t), "rxrings"},  

352     { MAC_PROP_TXRINGS RANGE, sizeof (mac_propval_range_t), "txrings"},  

354     { MAC_PROP_MAX_TX_RINGS_AVAIL, sizeof (uint_t), "txrings-available"},  

355  

357     { MAC_PROP_MAX_RX_RINGS_AVAIL, sizeof (uint_t), "rxrings-available"},  

358  

360     { MAC_PROP_MAX_RXHWCLNT_AVAIL, sizeof (uint_t), "rxhwclnt-available"},  

362     { MAC_PROP_MAX_TXHWCLNT_AVAIL, sizeof (uint_t), "txhwclnt-available"},  

364     { MAC_PROP_IB_LINKMODE,  sizeof (uint32_t),     "linkmode"},  

366     { MAC_PROP_MACADDRESS,  sizeof (mac_addrprop_t), "mac-address"},  

368     #endif /* ! codereview */  

369     { MAC_PROP_PRIVATE,      0,                      "driver-private"}  

370 };  

372 typedef struct bridge_public_prop_s {  

373     const char *bpp_name;  

374     int bpp_code;  

375 } bridge_public_prop_t;  

377 static const bridge_public_prop_t bridge_prop[] = {  

378     { "stp", PT_CFG_NON_STP },  

379     { "stp_priority", PT_CFG_PRIO },  

380     { "stp_cost", PT_CFG_COST },  

381     { "stp_edge", PT_CFG_EDGE },  

382     { "stp_p2p", PT_CFG_P2P }  


```

```

383     { "stp_mccheck", PT_CFG_MCHECK },
384     { NULL, 0 }
385 };
386
387 static val_desc_t link_duplex_vals[] = {
388     { "half", LINK_DUPLEX_HALF },
389     { "full", LINK_DUPLEX_HALF },
390 };
391 static val_desc_t link_status_vals[] = {
392     { "up", LINK_STATE_UP },
393     { "down", LINK_STATE_DOWN },
394 };
395 static val_desc_t link_01_vals[] = {
396     { "1", 1 },
397     { "0", 0 },
398 };
399 static val_desc_t link_flow_vals[] = {
400     { "no", LINK_FLOWCTRL_NONE },
401     { "tx", LINK_FLOWCTRL_TX },
402     { "rx", LINK_FLOWCTRL_RX },
403     { "bi", LINK_FLOWCTRL_BI },
404 };
405 static val_desc_t link_priority_vals[] = {
406     { "low", MPL_LOW },
407     { "medium", MPL_MEDIUM },
408     { "high", MPL_HIGH },
409 };
410
411 static val_desc_t link_tagmode_vals[] = {
412     { "normal", LINK_TAGMODE_NORMAL },
413     { "vlanonly", LINK_TAGMODE_VLANONLY },
414 };
415
416 static val_desc_t link_protect_vals[] = {
417     { "mac-nospoof", MPT_MACNOSPOOF },
418     { "restricted", MPT_RESTRICTED },
419     { "ip-nospoof", MPT_IPNOSPOOF },
420     { "dhcp-nospoof", MPT_DHCPNOSPOOF },
421 };
422
423 static val_desc_t dladm_wlan_radio_vals[] = {
424     { "on", DLADM_WLAN_RADIO_ON },
425     { "off", DLADM_WLAN_RADIO_OFF },
426 };
427
428 static val_desc_t dladm_wlan_powermode_vals[] = {
429     { "off", DLADM_WLAN_PM_OFF },
430     { "fast", DLADM_WLAN_PM_FAST },
431     { "max", DLADM_WLAN_PM_MAX },
432 };
433
434 static val_desc_t stp_p2p_vals[] = {
435     { "true", P2P_FORCE_TRUE },
436     { "false", P2P_FORCE_FALSE },
437     { "auto", P2P_AUTO },
438 };
439
440 static val_desc_t dladm_part_linkmode_vals[] = {
441     { "cm", DLADM_PART_CM_MODE },
442     { "ud", DLADM_PART_UD_MODE },
443 };
444
445 #define VALCNT(vals) (sizeof ((vals)) / sizeof (val_desc_t))
446 #define RESET_VAL ((uintptr_t)-1)
447 #define UNSPEC_VAL ((uintptr_t)-2)

```

```

448 static prop_desc_t prop_table[] = {
449     { "channel", { NULL, 0 },
450         NULL, 0, NULL, NULL,
451         get_channel, NULL, 0,
452         DATALINK_CLASS_PHYS, DL_WIFI },
453
454     { "powermode", { "off", DLADM_WLAN_PM_OFF },
455         dladm_wlan_powermode_vals, VALCNT(dladm_wlan_powermode_vals),
456         set_powermode, NULL,
457         get_powermode, NULL, 0,
458         DATALINK_CLASS_PHYS, DL_WIFI },
459
460     { "radio", { "on", DLADM_WLAN_RADIO_ON },
461         dladm_wlan_radio_vals, VALCNT(dladm_wlan_radio_vals),
462         set_radio, NULL,
463         get_radio, NULL, 0,
464         DATALINK_CLASS_PHYS, DL_WIFI },
465
466     { "linkmode", { "cm", DLADM_PART_CM_MODE },
467         dladm_part_linkmode_vals, VALCNT(dladm_part_linkmode_vals),
468         set_public_prop, NULL, get_linkmode_prop, NULL, 0,
469         DATALINK_CLASS_PART, DL_IB },
470
471     { "speed", { "", 0 }, NULL, 0,
472         set_rate, get_rate_mod,
473         get_rate, check_rate, 0,
474         DATALINK_CLASS_PHYS, DATALINK_ANY_MEDIATYPE },
475
476     { "autopush", { "", 0 }, NULL, 0,
477         set_public_prop, NULL,
478         get_autopush, check_autopush, PD_CHECK_ALLOC,
479         DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
480
481     { "zone", { "", 0 }, NULL, 0,
482         set_zone, NULL,
483         get_zone, check_zone, PD_TEMPONLY|PD_CHECK_ALLOC,
484         DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
485
486     { "duplex", { "", 0 },
487         link_duplex_vals, VALCNT(link_duplex_vals),
488         NULL, NULL, get_duplex, NULL,
489         0, DATALINK_CLASS_PHYS, DL_ETHER },
490
491     { "state", { "up", LINK_STATE_UP },
492         link_status_vals, VALCNT(link_status_vals),
493         NULL, NULL, get_link_state, NULL,
494         0, DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
495
496     { "adv_autoneg_cap", { "", 0 },
497         link_01_vals, VALCNT(link_01_vals),
498         set_public_prop, NULL, get_binary, NULL,
499         0, DATALINK_CLASS_PHYS, DL_ETHER },
500
501     { "mtu", { "", 0 }, NULL, 0,
502         set_public_prop, get_range,
503         get_uint32, check_uint32, 0, DATALINK_CLASS_ALL,
504         DATALINK_ANY_MEDIATYPE },
505
506     { "flowctrl", { "", 0 },
507         link_flow_vals, VALCNT(link_flow_vals),
508         set_public_prop, NULL, get_flowctl, NULL,
509         0, DATALINK_CLASS_PHYS, DL_ETHER },
510
511     { "adv_10gfidx_cap", { "", 0 },
512         link_01_vals, VALCNT(link_01_vals),
513         NULL, NULL, get_binary, NULL,
514 }

```

```

515      0, DATALINK_CLASS_PHYS, DL_ETHER },
516
517 { "en_10gfdx_cap", { "", 0 },
518   link_01_vals, VALCNT(link_01_vals),
519   set_public_prop, NULL, get_binary, NULL,
520   0, DATALINK_CLASS_PHYS, DL_ETHER },
521
522 { "adv_1000fdx_cap", { "", 0 },
523   link_01_vals, VALCNT(link_01_vals),
524   NULL, NULL, get_binary, NULL,
525   0, DATALINK_CLASS_PHYS, DL_ETHER },
526
527 { "en_1000fdx_cap", { "", 0 },
528   link_01_vals, VALCNT(link_01_vals),
529   set_public_prop, NULL, get_binary, NULL,
530   0, DATALINK_CLASS_PHYS, DL_ETHER },
531
532 { "adv_1000hdx_cap", { "", 0 },
533   link_01_vals, VALCNT(link_01_vals),
534   NULL, NULL, get_binary, NULL,
535   0, DATALINK_CLASS_PHYS, DL_ETHER },
536
537 { "en_1000hdx_cap", { "", 0 },
538   link_01_vals, VALCNT(link_01_vals),
539   set_public_prop, NULL, get_binary, NULL,
540   0, DATALINK_CLASS_PHYS, DL_ETHER },
541
542 { "adv_100fdx_cap", { "", 0 },
543   link_01_vals, VALCNT(link_01_vals),
544   NULL, NULL, get_binary, NULL,
545   0, DATALINK_CLASS_PHYS, DL_ETHER },
546
547 { "en_100fdx_cap", { "", 0 },
548   link_01_vals, VALCNT(link_01_vals),
549   set_public_prop, NULL, get_binary, NULL,
550   0, DATALINK_CLASS_PHYS, DL_ETHER },
551
552 { "adv_100hdx_cap", { "", 0 },
553   link_01_vals, VALCNT(link_01_vals),
554   NULL, NULL, get_binary, NULL,
555   0, DATALINK_CLASS_PHYS, DL_ETHER },
556
557 { "en_100hdx_cap", { "", 0 },
558   link_01_vals, VALCNT(link_01_vals),
559   set_public_prop, NULL, get_binary, NULL,
560   0, DATALINK_CLASS_PHYS, DL_ETHER },
561
562 { "adv_10fdx_cap", { "", 0 },
563   link_01_vals, VALCNT(link_01_vals),
564   NULL, NULL, get_binary, NULL,
565   0, DATALINK_CLASS_PHYS, DL_ETHER },
566
567 { "en_10fdx_cap", { "", 0 },
568   link_01_vals, VALCNT(link_01_vals),
569   set_public_prop, NULL, get_binary, NULL,
570   0, DATALINK_CLASS_PHYS, DL_ETHER },
571
572 { "adv_10hdx_cap", { "", 0 },
573   link_01_vals, VALCNT(link_01_vals),
574   NULL, NULL, get_binary, NULL,
575   0, DATALINK_CLASS_PHYS, DL_ETHER },
576
577 { "en_10hdx_cap", { "", 0 },
578   link_01_vals, VALCNT(link_01_vals),
579   set_public_prop, NULL, get_binary, NULL,
580   0, DATALINK_CLASS_PHYS, DL_ETHER },

```

```

582 { "maxbw", { "--", RESET_VAL }, NULL, 0,
583   set_resource, NULL,
584   get_maxbw, check_maxbw, PD_CHECK_ALLOC,
585   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
586
587 { "cpus", { "--", RESET_VAL }, NULL, 0,
588   set_resource, NULL,
589   get_cpus, check_cpus, 0,
590   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
591
592 { "cpus-effective", { "--", 0 },
593   NULL, 0, NULL, NULL,
594   get_cpus, 0, 0,
595   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
596
597 { "pool", { "--", RESET_VAL }, NULL, 0,
598   set_resource, NULL,
599   get_pool, check_pool, 0,
600   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
601
602 { "pool-effective", { "--", 0 },
603   NULL, 0, NULL, NULL,
604   get_pool, 0, 0,
605   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
606
607 { "priority", { "high", MPL_RESET },
608   link_priority_vals, VALCNT(link_priority_vals), set_resource,
609   NULL, get_priority, check_prop, 0,
610   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },
611
612 { "tagmode", { "vlanonly", LINK_TAGMODE_VLANONLY },
613   link_tagmode_vals, VALCNT(link_tagmode_vals),
614   set_public_prop, NULL, get_tagmode,
615   NULL, 0,
616   DATALINK_CLASS_PHYS | DATALINK_CLASS_AGGR | DATALINK_CLASS_VNIC,
617   DL_ETHER },
618
619 { "hoplimit", { "", 0 }, NULL, 0,
620   set_public_prop, get_range, get_uint32,
621   check_hoplimit, 0, DATALINK_CLASS_IPTUN, DATALINK_ANY_MEDIATYPE },
622
623 { "encaplimit", { "", 0 }, NULL, 0,
624   set_public_prop, get_range, get_uint32,
625   check_encaplim, 0, DATALINK_CLASS_IPTUN, DL_IPV6 },
626
627 { "forward", { "1", 1 },
628   link_01_vals, VALCNT(link_01_vals),
629   set_bridge_forward, NULL, get_bridge_forward, NULL, PD_AFTER_PERM,
630   DATALINK_CLASS_ALL & ~DATALINK_CLASS_VNIC, DL_ETHER },
631
632 { "default_tag", { "1", 1 }, NULL, 0,
633   set_bridge_pvid, NULL, get_bridge_pvid, check_bridge_pvid,
634   0, DATALINK_CLASS_PHYS | DATALINK_CLASS_AGGR |
635   DATALINK_CLASS_ETHERSTUB | DATALINK_CLASS_SIMNET, DL_ETHER },
636
637 { "learn_limit", { "1000", 1000 }, NULL, 0,
638   set_public_prop, NULL, get_uint32,
639   check_uint32, 0,
640   DATALINK_CLASS_PHYS | DATALINK_CLASS_AGGR |
641   DATALINK_CLASS_ETHERSTUB | DATALINK_CLASS_SIMNET, DL_ETHER },
642
643 { "learn_decay", { "200", 200 }, NULL, 0,
644   set_public_prop, NULL, get_uint32,
645   check_uint32, 0,
646   DATALINK_CLASS_PHYS | DATALINK_CLASS_AGGR |

```

new/usr/src/lib/libdladm/common/linkprop.c

9

```

447     DATALINK_CLASS_ETHERSTUB|DATALINK_CLASS_SIMNET, DL_ETHER },

448 { "stp", { "1", 1 },
449   link_01_vals, VALCNT(link_01_vals),
450   set_stp_prop, NULL, get_stp, NULL, PD_AFTER_PERM,
451   DATALINK_CLASS_PHYS|DATALINK_CLASS_AGGR|
452   DATALINK_CLASS_ETHERSTUB|DATALINK_CLASS_SIMNET, DL_ETHER },

453 { "stp_priority", { "128", 128 }, NULL, 0,
454   set_stp_prop, NULL, get_stp, check_stp_prop, PD_AFTER_PERM,
455   DATALINK_CLASS_PHYS|DATALINK_CLASS_AGGR|
456   DATALINK_CLASS_ETHERSTUB|DATALINK_CLASS_SIMNET, DL_ETHER },

457 { "stp_cost", { "auto", 0 }, NULL, 0,
458   set_stp_prop, NULL, get_stp, check_stp_prop, PD_AFTER_PERM,
459   DATALINK_CLASS_PHYS|DATALINK_CLASS_AGGR|
460   DATALINK_CLASS_ETHERSTUB|DATALINK_CLASS_SIMNET, DL_ETHER },

461 { "stp_edge", { "1", 1 },
462   link_01_vals, VALCNT(link_01_vals),
463   set_stp_prop, NULL, get_stp, NULL, PD_AFTER_PERM,
464   DATALINK_CLASS_PHYS|DATALINK_CLASS_AGGR|
465   DATALINK_CLASS_ETHERSTUB|DATALINK_CLASS_SIMNET, DL_ETHER },

466 { "stp_p2p", { "auto", P2P_AUTO },
467   stp_p2p_vals, VALCNT(stp_p2p_vals),
468   set_stp_prop, NULL, get_stp, NULL, PD_AFTER_PERM,
469   DATALINK_CLASS_PHYS|DATALINK_CLASS_AGGR|
470   DATALINK_CLASS_ETHERSTUB|DATALINK_CLASS_SIMNET, DL_ETHER },

471 { "stp_mcheck", { "0", 0 },
472   link_01_vals, VALCNT(link_01_vals),
473   set_stp_prop, NULL, get_stp, check_stp_prop, PD_AFTER_PERM,
474   DATALINK_CLASS_PHYS|DATALINK_CLASS_AGGR|
475   DATALINK_CLASS_ETHERSTUB|DATALINK_CLASS_SIMNET, DL_ETHER },

476 { "protection", { "--", RESET_VAL },
477   link_protect_vals, VALCNT(link_protect_vals),
478   set_resource, NULL, get_protection, check_prop, 0,
479   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },

480 { "mac-address", { "", 0 }, NULL, 0,
481   set_macaddr, NULL, get_macaddr, check_macaddr, PD_CHECK_ALLOC,
482   DATALINK_CLASS_PHYS|DATALINK_CLASS_AGGR|DATALINK_CLASS_VNIC|
483   DATALINK_CLASS_SIMNET, DATALINK_ANY_MEDIATYPE },

484 #endif /* ! codereview */
485 { "allowed-ips", { "--", 0 },
486   NULL, 0, set_resource, NULL,
487   get_allowedips, check_allowedips, PD_CHECK_ALLOC,
488   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },

489 { "allowed-dhcp-cids", { "--", 0 },
490   NULL, 0, set_resource, NULL,
491   get_allowedcids, check_allowedcids, PD_CHECK_ALLOC,
492   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },

493 { "rxrings", { "--", RESET_VAL }, NULL, 0,
494   set_resource, get_rings_range, get_rxrings, check_rings, 0,
495   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },

496 { "rxrings-effective", { "--", 0 },
497   NULL, 0, NULL, NULL,
498   get_rxrings, NULL, 0,
499   DATALINK_CLASS_ALL, DATALINK_ANY_MEDIATYPE },

```

new/usr/src/lib/libdladm/common/linkprop.c

```

2292     datalink_media_t media)
2293 {
2294     mac_addrprop_t *addrprop;
2295     val_desc_t      *vdp = *vdpp;
2296     uint_t          val_cnt = *val_cntp;
2297     uchar_t         *macaddr;
2298     int             maclen;
2299     dladm_status_t  status = DLADM_STATUS_OK;
2300
2301     if (val_cnt != 1)
2302         return (DLADM_STATUS_BADVALCNT);
2303
2304     macaddr = _link_aton(*prop_val, &maclen);
2305     if (macaddr == NULL)
2306         return (DLADM_STATUS_BADVAL);
2307
2308     addrprop = malloc(sizeof (mac_addrprop_t));
2309     if (addrprop == NULL) {
2310         status = DLADM_STATUS_NOMEM;
2311         goto out;
2312     }
2313
2314     (void) memcpy(addrprop->ma_addr, macaddr, maclen);
2315     addrprop->ma_len = maclen;
2316
2317     vdp->vd_val = (uintptr_t)addrprop;
2318
2319 out:
2320     free(macaddr);
2321     return (status);
2322 }
2323 /* ARGSUSED */
2324 static dladm_status_t
2325 #endif /* ! codereview */
2326 set_resource(dladm_handle_t handle, prop_desc_t *pdp,
2327               datalink_id_t linkid, val_desc_t *vdp, uint_t val_cnt,
2328               uint_t flags, datalink_media_t media)
2329 {
2330     mac_resource_props_t    mrp;
2331     dladm_status_t          status = DLADM_STATUS_OK;
2332     dld_ioc_macprop_t       dip;
2333     int                      i;
2334
2335     bzero(&mrp, sizeof (mac_resource_props_t));
2336     dip = i_dladm_buf_alloc_by_name(0, linkid, "resource",
2337                                     flags, &status);
2338
2339     if (dip == NULL)
2340         return (status);
2341
2342     for (i = 0; i < DLADM_MAX_RSRC_PROP; i++) {
2343         resource_prop_t *rp = &rsrc_prop_table[i];
2344
2345         if (strcmp(pdp->pd_name, rp->rp_name) != 0)
2346             continue;
2347
2348         status = rp->rp_extract(vdp, val_cnt, &mrp);
2349         if (status != DLADM_STATUS_OK)
2350             goto done;
2351
2352         break;
2353     }
2354
2355     (void) memcpy(dip->pr_val, &mrp, dip->pr_valsize);
2356     status = i_dladm_macprop(handle, dip, B_TRUE);
2357

```

```

2358 done:
2359     free(dip);
2360     return (status);
2361 }
2362
2363 /* ARGSUSED */
2364 static dladm_status_t
2365 get_protection(dladm_handle_t handle, prop_desc_t *pdp,
2366                 datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
2367                 datalink_media_t media, uint_t flags, uint_t *perm_flags)
2368 {
2369     mac_resource_props_t    mrp;
2370     mac_protect_t           *p;
2371     dladm_status_t          status;
2372     uint32_t                i, cnt = 0, setbits[32];
2373
2374     status = i_dladm_get_public_prop(handle, linkid, "resource", flags,
2375                                     perm_flags, &mrp, sizeof (mrp));
2376     if (status != DLADM_STATUS_OK)
2377         return (status);
2378
2379     p = &mrp.mrp_protect;
2380     if ((mrp.mrp_mask & MRP_PROTECT) == 0) {
2381         *val_cnt = 0;
2382         return (DLADM_STATUS_OK);
2383     }
2384     dladm_find_setbits32(p->mp_types, setbits, &cnt);
2385     if (cnt > *val_cnt)
2386         return (DLADM_STATUS_BADVALCNT);
2387
2388     for (i = 0; i < cnt; i++)
2389         (void) dladm_protect2str(setbits[i], prop_val[i]);
2390
2391     *val_cnt = cnt;
2392     return (DLADM_STATUS_OK);
2393 }
2394
2395 /* ARGSUSED */
2396 static dladm_status_t
2397 get_allowedips(dladm_handle_t handle, prop_desc_t *pdp,
2398                  datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
2399                  datalink_media_t media, uint_t flags, uint_t *perm_flags)
2400 {
2401     mac_resource_props_t    mrp;
2402     mac_protect_t           *p;
2403     dladm_status_t          status;
2404     int                      i;
2405
2406     status = i_dladm_get_public_prop(handle, linkid, "resource", flags,
2407                                     perm_flags, &mrp, sizeof (mrp));
2408     if (status != DLADM_STATUS_OK)
2409         return (status);
2410
2411     p = &mrp.mrp_protect;
2412     if (p->mp_ipaddrcnt == 0) {
2413         *val_cnt = 0;
2414         return (DLADM_STATUS_OK);
2415     }
2416     if (p->mp_ipaddrcnt > *val_cnt)
2417         return (DLADM_STATUS_BADVALCNT);
2418
2419     for (i = 0; i < p->mp_ipaddrcnt; i++) {
2420         if (p->mp_ipaddrs[i].ip_version == IPV4_VERSION) {
2421             ipaddr_t              v4addr;
2422

```

```

2424             v4addr = V4_PART_OF_V6(p->mp_ipaddrs[i].ip_addr);
2425             (void) dladm_ipv4addr2str(&v4addr, prop_val[i]);
2426         } else {
2427             (void) dladm_ipv6addr2str(&p->mp_ipaddrs[i].ip_addr,
2428                                       prop_val[i]);
2429         }
2430     }
2431     *val_cnt = p->mp_ipaddrcnt;
2432     return (DLADM_STATUS_OK);
2433 }

2435 dladm_status_t
2436 extract_protection(val_desc_t *vdpp, uint_t cnt, void *arg)
2437 {
2438     mac_resource_props_t    *mrp = arg;
2439     uint32_t                 types = 0;
2440     int                      i;
2441
2442     for (i = 0; i < cnt; i++)
2443         types |= (uint32_t)vdpp[i].vd_val;
2444
2445     mrp->mrp_protect.mp_types = types;
2446     mrp->mrp_mask |= MRP_PROTECT;
2447     return (DLADM_STATUS_OK);
2448 }

2450 dladm_status_t
2451 extract_allowedips(val_desc_t *vdpp, uint_t cnt, void *arg)
2452 {
2453     mac_resource_props_t    *mrp = arg;
2454     mac_protect_t           *p = &mrp->mrp_protect;
2455     int                      i;
2456
2457     if (vdpp->vd_val == 0) {
2458         cnt = (uint_t)-1;
2459     } else {
2460         for (i = 0; i < cnt; i++) {
2461             bcopy((void *)vdpp[i].vd_val, &p->mp_ipaddrs[i],
2462                   sizeof (mac_ipaddr_t));
2463         }
2464     }
2465     p->mp_ipaddrcnt = cnt;
2466     mrp->mrp_mask |= MRP_PROTECT;
2467     return (DLADM_STATUS_OK);
2468 }

2470 static dladm_status_t
2471 check_single_ip(char *buf, mac_ipaddr_t *addr)
2472 {
2473     dladm_status_t status;
2474     ipaddr_t      v4addr;
2475     in6_addr_t    v6addr;
2476     boolean_t      isv4 = B_TRUE;
2477
2478     status = dladm_str2ipv4addr(buf, &v4addr);
2479     if (status == DLADM_STATUS_INVALID_IP) {
2480         status = dladm_str2ipv6addr(buf, &v6addr);
2481         if (status == DLADM_STATUS_OK)
2482             isv4 = B_FALSE;
2483     }
2484     if (status != DLADM_STATUS_OK)
2485         return (status);
2486
2487     if (isv4) {
2488         if (v4addr == INADDR_ANY)
2489             return (DLADM_STATUS_INVALID_IP);

```

```

2491             IN6_IPADDR_TO_V4MAPPED(v4addr, &addr->ip_addr);
2492             addr->ip_version = IPV4_VERSION;
2493         } else {
2494             if (IN6_IS_ADDR_UNSPECIFIED(&v6addr))
2495                 return (DLADM_STATUS_INVALID_IP);
2496
2497             addr->ip_addr = v6addr;
2498             addr->ip_version = IPV6_VERSION;
2499         }
2500     }
2501     return (DLADM_STATUS_OK);

2503 /* ARGSUSED */
2504 static dladm_status_t
2505 check_allowedips(dladm_handle_t handle, prop_desc_t *pdp,
2506                    datalink_id_t linkid, char **prop_val, uint_t *val_cntp, uint_t flags,
2507                    val_desc_t **vdpp, datalink_media_t media)
2508 {
2509     dladm_status_t status;
2510     mac_ipaddr_t   *addr;
2511     int              i;
2512     uint_t           val_cnt = *val_cntp;
2513     val_desc_t      *vdpp = *vdpp;
2514
2515     if (val_cnt > MPT_MAXIPADDR)
2516         return (DLADM_STATUS_BADVALCNT);
2517
2518     for (i = 0; i < val_cnt; i++) {
2519         if ((addr = calloc(1, sizeof (mac_ipaddr_t))) == NULL) {
2520             status = DLADM_STATUS_NOMEM;
2521             goto fail;
2522         }
2523         vdpp[i].vd_val = (uintptr_t)addr;
2524
2525         status = check_single_ip(prop_val[i], addr);
2526         if (status != DLADM_STATUS_OK)
2527             goto fail;
2528     }
2529     return (DLADM_STATUS_OK);

2531 fail:
2532     for (i = 0; i < val_cnt; i++) {
2533         free((void *)vdpp[i].vd_val);
2534         vdpp[i].vd_val = NULL;
2535     }
2536     return (status);

2539 static void
2540 dladm_cid2str(mac_dhcpcid_t *cid, char *buf)
2541 {
2542     char tmp_buf[DLADM_STRSIZE];
2543     uint_t hexlen;
2544
2545     switch (cid->dc_form) {
2546     case CIDFORM_TYPED: {
2547         uint16_t      duidtype, hwtype;
2548         uint32_t      timestamp, ennum;
2549         char          *lladdr;
2550
2551         if (cid->dc_len < sizeof (duidtype))
2552             goto fail;
2553
2554         bcopy(cid->dc_id, &duidtype, sizeof (duidtype));
2555         duidtype = ntohs(duidtype);
2556     }
2557 }

```

```

2556         switch (duidtype) {
2557             case DHCPV6_DUID_LL:
2558                 duid_llt_t llt;
2559
2560                 if (cid->dc_len < sizeof (llt))
2561                     goto fail;
2562
2563                 bcopy(cid->dc_id, &llt, sizeof (llt));
2564                 hwtype = ntohs(llt.dllt_hwtype);
2565                 timestamp = ntohl(llt.dllt_time);
2566                 lladdr = _link_ntoa(cid->dc_id + sizeof (llt),
2567                                     NULL, cid->dc_len - sizeof (llt), IFT_OTHER);
2568                 if (lladdr == NULL)
2569                     goto fail;
2570
2571                 (void) sprintf(buf, DLADM_STRSIZE, "%d.%d.%d.%s",
2572                               duidtype, hwtype, timestamp, lladdr);
2573                 free(lladdr);
2574                 break;
2575
2576             case DHCPV6_DUID_EN:
2577                 duid_en_t en;
2578
2579                 if (cid->dc_len < sizeof (en))
2580                     goto fail;
2581
2582                 bcopy(cid->dc_id, &en, sizeof (en));
2583                 ennum = DHCPV6_GET_ENTNUM(&en);
2584                 hexlen = sizeof (tmp_buf);
2585                 if (octet_to_hexascii(cid->dc_id + sizeof (en),
2586                                       cid->dc_len - sizeof (en), tmp_buf, &hexlen) != 0)
2587                     goto fail;
2588
2589                 (void) sprintf(buf, DLADM_STRSIZE, "%d.%d.%s",
2590                               duidtype, ennum, tmp_buf);
2591                 break;
2592
2593             case DHCPV6_DUID_LL:
2594                 duid_llt_t ll;
2595
2596                 if (cid->dc_len < sizeof (ll))
2597                     goto fail;
2598
2599                 bcopy(cid->dc_id, &ll, sizeof (ll));
2600                 hwtype = ntohs(ll.dllt_hwtype);
2601                 lladdr = _link_ntoa(ll.dllt_dc_id + sizeof (ll),
2602                                     NULL, cid->dc_len - sizeof (ll), IFT_OTHER);
2603                 if (lladdr == NULL)
2604                     goto fail;
2605
2606                 (void) sprintf(buf, DLADM_STRSIZE, "%d.%d.%s",
2607                               duidtype, hwtype, lladdr);
2608                 free(lladdr);
2609                 break;
2610
2611         default:
2612             hexlen = sizeof (tmp_buf);
2613             if (octet_to_hexascii(cid->dc_id + sizeof (duidtype),
2614                                   cid->dc_len - sizeof (duidtype),
2615                                   tmp_buf, &hexlen) != 0)
2616                 goto fail;
2617
2618             (void) sprintf(buf, DLADM_STRSIZE, "%d.%s",
2619                           duidtype, tmp_buf);
2620
2621     }

```

```

2622             break;
2623
2624         case CIDFORM_HEX:
2625             hexlen = sizeof (tmp_buf);
2626             if (octet_to_hexascii(cid->dc_id, cid->dc_len,
2627                                   tmp_buf, &hexlen) != 0)
2628                 goto fail;
2629
2630             (void) sprintf(buf, DLADM_STRSIZE, "0x%s", tmp_buf);
2631             break;
2632
2633         case CIDFORM_STR:
2634             int i;
2635
2636             for (i = 0; i < cid->dc_len; i++) {
2637                 if (!isprint(cid->dc_id[i]))
2638                     goto fail;
2639             }
2640             (void) sprintf(buf, DLADM_STRSIZE, "%s", cid->dc_id);
2641             break;
2642
2643         default:
2644             goto fail;
2645
2646     return;
2647
2648 fail:
2649     (void) sprintf(buf, DLADM_STRSIZE, "<unknown>");
2650 }
2651
2652 static dladm_status_t
2653 dladm_str2cid(char *buf, mac_dhcpcid_t *cid)
2654 {
2655     char *ptr = buf;
2656     char tmp_buf[DLADM_STRSIZE];
2657     uint_t hexlen, cidlen;
2658
2659     bzero(cid, sizeof (*cid));
2660     if (isdigit(*ptr) &&
2661         ptr[strspn(ptr, "0123456789")] == '.') {
2662         char *cp;
2663         ulong_t duidtype;
2664         ulong_t subtype;
2665         ulong_t timestamp;
2666         uchar_t *lladdr;
2667         int addrlen;
2668
2669         errno = 0;
2670         duidtype = strtoul(ptr, &cp, 0);
2671         if (ptr == cp || errno != 0 || *cp != '.' || duidtype > USHRT_MAX)
2672             return (DLADM_STATUS_BADARG);
2673         ptr = cp + 1;
2674
2675         if (duidtype != 0 && duidtype <= DHCPV6_DUID_LL) {
2676             errno = 0;
2677             subtype = strtoul(ptr, &cp, 0);
2678             if (ptr == cp || errno != 0 || *cp != '.')
2679                 return (DLADM_STATUS_BADARG);
2680             ptr = cp + 1;
2681
2682         }
2683         switch (duidtype) {
2684             case DHCPV6_DUID_LL:
2685                 duid_llt_t llt;
2686
2687                 errno = 0;

```

```

2688     timestamp = strtoul(ptr, &cp, 0);
2689     if (ptr == cp || errno != 0 || *cp != '.')
2690         return (DLADM_STATUS_BADARG);

2692     ptr = cp + 1;
2693     lladdr = _link_aton(ptr, &addrlen);
2694     if (lladdr == NULL)
2695         return (DLADM_STATUS_BADARG);

2697     cidlen = sizeof (llt) + addrlen;
2698     if (cidlen > sizeof (cid->dc_id)) {
2699         free(lladdr);
2700         return (DLADM_STATUS_TOOSMALL);
2701     }
2702     llt.dllt_dutype = htons(duidtype);
2703     llt.dllt_hwtype = htons(subtype);
2704     llt.dllt_time = htonl(timestamp);
2705     bcopy(&llt, cid->dc_id, sizeof (llt));
2706     bcopy(lladdr, cid->dc_id + sizeof (llt), addrlen);
2707     free(lladdr);
2708     break;
2709 }
2710 case DHCPV6_DUID_LL: {
2711     duid_ll_t      ll;
2712
2713     lladdr = _link_aton(ptr, &addrlen);
2714     if (lladdr == NULL)
2715         return (DLADM_STATUS_BADARG);

2717     cidlen = sizeof (ll) + addrlen;
2718     if (cidlen > sizeof (cid->dc_id)) {
2719         free(lladdr);
2720         return (DLADM_STATUS_TOOSMALL);
2721     }
2722     ll.dll_dutype = htons(duidtype);
2723     ll.dll_hwtype = htons(subtype);
2724     bcopy(&ll, cid->dc_id, sizeof (ll));
2725     bcopy(lladdr, cid->dc_id + sizeof (ll), addrlen);
2726     free(lladdr);
2727     break;
2728 }
2729 default: {
2730     hexlen = sizeof (tmp_buf);
2731     if (hexascii_to_octet(ptr, strlen(ptr),
2732                           tmp_buf, &hexlen) != 0)
2733         return (DLADM_STATUS_BADARG);

2735     if (duidtype == DHCPV6_DUID_EN) {
2736         duid_en_t      en;

2738         en.den_dutype = htons(duidtype);
2739         DHCPV6_SET_ENTNUM(&en, subtype);

2741         cidlen = sizeof (en) + hexlen;
2742         if (cidlen > sizeof (cid->dc_id))
2743             return (DLADM_STATUS_TOOSMALL);

2745         bcopy(&en, cid->dc_id, sizeof (en));
2746         bcopy(tmp_buf, cid->dc_id + sizeof (en),
2747               hexlen);
2748     } else {
2749         uint16_t        dutype = htons(duidtype);

2751         cidlen = sizeof (dutype) + hexlen;
2752         if (cidlen > sizeof (cid->dc_id))
2753             return (DLADM_STATUS_TOOSMALL);

```

```

2755
2756
2757
2758
2759
2760
2761
2762
2763
2764
2765
2766
2767
2768
2769
2770
2771
2772
2773
2774
2775
2776
2777
2778
2779
2780
2781
2782
2783
2784
2785
2786
2787
2788 /* ARGUSED */
2789 static dladm_status_t
2790 get_allowedcids(dladm_handle_t handle, prop_desc_t *pdp,
2791                  datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
2792                  datalink_media_t media, uint_t flags, uint_t *perm_flags)
2793 {
2794     mac_resource_props_t    mrp;
2795     mac_protect_t          *p;
2796     dladm_status_t          status;
2797     int                      i;

2798     status = i_dladm_get_public_prop(handle, linkid, "resource", flags,
2799                                     perm_flags, &mrp, sizeof (mrp));
2800     if (status != DLADM_STATUS_OK)
2801         return (status);

2802     p = &mrp.mrp_protect;
2803     if (p->mp_cidcnt == 0) {
2804         *val_cnt = 0;
2805         return (DLADM_STATUS_OK);
2806     }
2807     if (p->mp_cidcnt > *val_cnt)
2808         return (DLADM_STATUS_BADVALCNT);

2809     for (i = 0; i < p->mp_cidcnt; i++) {
2810         mac_dhcpcid_t *cid = &p->mp_cids[i];
2811
2812         dladm_cid2str(cid, prop_val[i]);
2813     }
2814     *val_cnt = p->mp_cidcnt;
2815
2816
2817
2818
2819 }

```

```

2821 dladm_status_t
2822 extract_allowedcids(val_desc_t *vdp, uint_t cnt, void *arg)
2823 {
2824     mac_resource_props_t      *mrp = arg;
2825     mac_protect_t            *p = &mrp->mrp_protect;
2826     int                      i;
2827
2828     if (vdp->vd_val == 0) {
2829         cnt = (uint_t)-1;
2830     } else {
2831         for (i = 0; i < cnt; i++) {
2832             bcopy((void *)vdp[i].vd_val, &p->mp_cids[i],
2833                   sizeof (mac_dhcpcid_t));
2834         }
2835     }
2836     p->mp_cidcnt = cnt;
2837     mrp->mrp_mask |= MRP_PROTECT;
2838     return (DLADM_STATUS_OK);
2839 }
2840
2841 /* ARGSUSED */
2842 static dladm_status_t
2843 check_allowedcids(dladm_handle_t handle, prop_desc_t *pdp,
2844     datalink_id_t linkid, char **prop_val, uint_t *val_cntp,
2845     uint_t flags, val_desc_t **vdpp, datalink_media_t media)
2846 {
2847     dladm_status_t    status;
2848     mac_dhcpcid_t   *cid;
2849     int              i;
2850     uint_t           val_cnt = *val_cntp;
2851     val_desc_t       *vdp = *vdpp;
2852
2853     if (val_cnt > MPT_MAXCID)
2854         return (DLADM_STATUS_BADVALCNT);
2855
2856     for (i = 0; i < val_cnt; i++) {
2857         if ((cid = calloc(1, sizeof (mac_dhcpcid_t))) == NULL) {
2858             status = DLADM_STATUS_NOMEM;
2859             goto fail;
2860         }
2861         vdp[i].vd_val = (uintptr_t)cid;
2862
2863         status = dladm_str2cid(prop_val[i], cid);
2864         if (status != DLADM_STATUS_OK)
2865             goto fail;
2866     }
2867     return (DLADM_STATUS_OK);
2868
2869 fail:
2870     for (i = 0; i < val_cnt; i++) {
2871         free((void *)vdp[i].vd_val);
2872         vdp[i].vd_val = NULL;
2873     }
2874     return (status);
2875 }
2876
2877 /* ARGSUSED */
2878 static dladm_status_t
2879 get_autopush(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
2880     char **prop_val, uint_t *val_cnt, datalink_media_t media,
2881     uint_t flags, uint_t *perm_flags)
2882 {
2883     struct          dlautopush dlap;
2884     int             i, len;
2885     dladm_status_t status;

```

```

2887     if (flags & DLD_PROP_DEFAULT)
2888         return (DLADM_STATUS_NOTDEFINED);
2889
2890     status = i_dladm_get_public_prop(handle, linkid, pdp->pd_name, flags,
2891                                     perm_flags, &dlap, sizeof(dlap));
2892     if (status != DLADM_STATUS_OK)
2893         return (status);
2894
2895     if (dlap.dap_npush == 0) {
2896         *val_cnt = 0;
2897         return (DLADM_STATUS_OK);
2898     }
2899     for (i = 0, len = 0; i < dlap.dap_npush; i++) {
2900         if (i != 0) {
2901             (void) sprintf(*prop_val + len,
2902                           DLADM_PROP_VAL_MAX - len, "%c",
2903                           AP_DELIMITER);
2904             len += 1;
2905         }
2906         (void) sprintf(*prop_val + len, DLADM_PROP_VAL_MAX - len,
2907                       "%s", dlap.dap_aplist[i]);
2908         len += strlen(dlap.dap_aplist[i]);
2909         if (dlap.dap_anchor - 1 == i) {
2910             (void) sprintf(*prop_val + len,
2911                           DLADM_PROP_VAL_MAX - len, "%c%s",
2912                           AP_DELIMITER,
2913                           AP_ANCHOR);
2914             len += (strlen(AP_ANCHOR) + 1);
2915         }
2916     }
2917     *val_cnt = 1;
2918     return (DLADM_STATUS_OK);
2919 }
2920 */
2921 * Add the specified module to the dlautopush structure; returns a
2922 * DLADM_STATUS_* code.
2923 */
2924 dladm_status_t
2925 i_dladm_add_ap_module(const char *module, struct dlautopush *dlap)
2926 {
2927     if ((strlen(module) == 0) || (strlen(module) > FMNAMESZ))
2928         return (DLADM_STATUS_BADVAL);
2929
2930     if (strncasecmp(module, AP_ANCHOR, strlen(AP_ANCHOR)) == 0) {
2931         /*
2932          * We don't allow multiple anchors, and the anchor must
2933          * be after at least one module.
2934          */
2935         if (dlap->dap_anchor != 0)
2936             return (DLADM_STATUS_BADVAL);
2937         if (dlap->dap_npush == 0)
2938             return (DLADM_STATUS_BADVAL);
2939
2940         dlap->dap_anchor = dlap->dap_npush;
2941         return (DLADM_STATUS_OK);
2942     }
2943     if (dlap->dap_npush >= MAXAPUSH)
2944         return (DLADM_STATUS_BADVALCNT);
2945
2946     (void) strlcpy(dlap->dap_aplist[dlap->dap_npush++], module,
2947                   FMNAMESZ + 1);
2948
2949     return (DLADM_STATUS_OK);
2950 }

```

```

2952 * Currently, both '.' and ' ' (space) can be used as the delimiters between
2953 * autopush modules. The former is used in dladm set-linkprop, and the
2954 * latter is used in the autopush(1M) file.
2955 */
2956 /* ARGSUSED */
2957 static dladm_status_t
2958 check_autopush(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
2959     char **prop_val, uint_t *val_cntp, uint_t flags, val_desc_t **vdpp,
2960     datalink_media_t media)
2961 {
2962     char *module;
2963     struct dlautopush *dlap;
2964     dladm_status_t status;
2965     char val[DLADM_PROP_VAL_MAX];
2966     char delimiter[4];
2967     uint_t val_cnt = *val_cntp;
2968     val_desc_t *vdp = *vdpp;
2969
2970     if (val_cnt != 1)
2971         return (DLADM_STATUS_BADVALCNT);
2972
2973     if (prop_val != NULL) {
2974         dlap = malloc(sizeof (struct dlautopush));
2975         if (dlap == NULL)
2976             return (DLADM_STATUS_NOMEM);
2977
2978         (void) memset(dlap, 0, sizeof (struct dlautopush));
2979         (void) snprintf(delimiter, 4, " %c\n", AP_DELIMITER);
2980         bcopy(*prop_val, val, DLADM_PROP_VAL_MAX);
2981         module = strtok(val, delimiter);
2982         while (module != NULL) {
2983             status = i_dladm_add_ap_module(module, dlap);
2984             if (status != DLADM_STATUS_OK)
2985                 return (status);
2986             module = strtok(NULL, delimiter);
2987         }
2988
2989         vdp->vd_val = (uintptr_t)dlap;
2990     } else {
2991         vdp->vd_val = 0;
2992     }
2993     return (DLADM_STATUS_OK);
2994 }
2995
2996 #define WLDP_BUFSIZE (MAX_BUF_LEN - WIFI_BUF_OFFSET)
2997
2998 /* ARGSUSED */
2999 static dladm_status_t
3000 get_rate_common(dladm_handle_t handle, prop_desc_t *pdp,
3001     datalink_id_t linkid, char **prop_val, uint_t *val_cntp, uint_t id,
3002     uint_t *perm_flags)
3003 {
3004     wl_rates_t *wrp;
3005     uint_t i;
3006     dladm_status_t status = DLADM_STATUS_OK;
3007
3008     wrp = malloc(WLDP_BUFSIZE);
3009     if (wrp == NULL)
3010         return (DLADM_STATUS_NOMEM);
3011
3012     status = i_dladm_wlan_param(handle, linkid, wrp, id, WLDP_BUFSIZE,
3013         B_FALSE);
3014     if (status != DLADM_STATUS_OK)
3015         goto done;
3016
3017     if (wrp->wl_rates_num > *val_cntp) {

```

```

3018         status = DLADM_STATUS_TOOSMALL;
3019         goto done;
3020     }
3021
3022     if (wrp->wl_rates_rates[0] == 0) {
3023         prop_val[0][0] = '\0';
3024         *val_cnt = 1;
3025         goto done;
3026     }
3027
3028     for (i = 0; i < wrp->wl_rates_num; i++) {
3029         (void) sprintf(prop_val[i], DLADM_STRSIZE, "%.*f",
3030             wrp->wl_rates_rates[i] % 2,
3031             (float)wrp->wl_rates_rates[i] / 2);
3032     }
3033     *val_cnt = wrp->wl_rates_num;
3034     *perm_flags = MAC_PROP_PERM_RW;
3035
3036 done:
3037     free(wrp);
3038     return (status);
3039 }
3040
3041 static dladm_status_t
3042 get_rate(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3043     char **prop_val, uint_t *val_cntp, datalink_media_t media,
3044     uint_t flags, uint_t *perm_flags)
3045 {
3046     if (media != DL_WIFI) {
3047         return (get_speed(handle, pdp, linkid, prop_val,
3048             val_cntp, media, flags, perm_flags));
3049     }
3050
3051     return (get_rate_common(handle, pdp, linkid, prop_val, val_cntp,
3052         MAC_PROP_WL_DESIRED_RATES, perm_flags));
3053 }
3054
3055 /* ARGSUSED */
3056 static dladm_status_t
3057 get_rate_mod(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3058     char **prop_val, uint_t *val_cntp, datalink_media_t media,
3059     uint_t flags, uint_t *perm_flags)
3060 {
3061     switch (media) {
3062     case DL_ETHER:
3063         /*
3064          * Speed for ethernet links is unbounded. E.g., 802.11b
3065          * links can have a speed of 5.5 Gbps.
3066         */
3067         return (DLADM_STATUS_NOTSUP);
3068
3069     case DL_WIFI:
3070         return (get_rate_common(handle, pdp, linkid, prop_val,
3071             val_cntp, MAC_PROP_WL_SUPPORTED_RATES, perm_flags));
3072     default:
3073         return (DLADM_STATUS_BADARG);
3074     }
3075 }
3076
3077 static dladm_status_t
3078 set_wlan_rate(dladm_handle_t handle, datalink_id_t linkid,
3079     dladm_wlan_rates_t *rates)
3080 {
3081     int i;
3082     uint_t len;
3083     wl_rates_t *wrp;

```

```

3084     dladm_status_t status = DLADM_STATUS_OK;
3085
3086     wrp = malloc(WLDP_BUFSIZE);
3087     if (wrp == NULL)
3088         return (DLADM_STATUS_NOMEM);
3089
3090     bzero(wrp, WLDP_BUFSIZE);
3091     for (i = 0; i < rates->wr_cnt; i++)
3092         wrp->wl_rates_rates[i] = rates->wr_rates[i];
3093     wrp->wl_rates_num = rates->wr_cnt;
3094
3095     len = offsetof(wl_rates_t, wl_rates_rates) +
3096           (rates->wr_cnt * sizeof (char)) + WIFI_BUF_OFFSET;
3097     status = i_dladm_wlan_param(handle, linkid, wrp,
3098                               MAC_PROP_WL_DESIRED_RATES, len, B_TRUE);
3099
3100     free(wrp);
3101     return (status);
3102 }
3103 /* ARGSUSED */
3104 static dladm_status_t
3105 set_rate(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3106           val_desc_t *vdp, uint_t val_cnt, uint_t flags, datalink_media_t media)
3107 {
3108     dladm_wlan_rates_t rates;
3109     dladm_status_t status;
3110
3111     /*
3112      * can currently set rate on WIFI links only.
3113      */
3114     if (media != DL_WIFI)
3115         return (DLADM_STATUS_PROPRDONLY);
3116
3117     if (val_cnt != 1)
3118         return (DLADM_STATUS_BADVALCNT);
3119
3120     rates.wr_cnt = 1;
3121     rates.wr_rates[0] = vdp[0].vd_val;
3122
3123     status = set_wlan_rate(handle, linkid, &rates);
3124
3125     return (status);
3126 }
3127 */
3128 /* ARGSUSED */
3129 static dladm_status_t
3130 check_rate(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3131             char **prop_val, uint_t *val_cntp, uint_t flags, val_desc_t **vdpp,
3132             datalink_media_t media)
3133 {
3134     int i;
3135     uint_t modval_cnt = MAX_SUPPORT_RATES;
3136     char *buf, *modval;
3137     dladm_status_t status;
3138     uint_t perm_flags;
3139     uint_t val_cnt = *val_cntp;
3140     val_desc_t *vdp = *vdpp;
3141
3142     if (val_cnt != 1)
3143         return (DLADM_STATUS_BADVALCNT);
3144
3145     buf = malloc(sizeof (char *) + DLADM_STRSIZE) *
3146           MAX_SUPPORT_RATES;
3147     if (buf == NULL) {
3148         status = DLADM_STATUS_NOMEM;
3149

```

```

3150             goto done;
3151     }
3152
3153     modval = (char **)(void *)buf;
3154     for (i = 0; i < MAX_SUPPORT_RATES; i++) {
3155         modval[i] = buf + sizeof (char *) * MAX_SUPPORT_RATES +
3156                     i * DLADM_STRSIZE;
3157     }
3158
3159     status = get_rate_mod(handle, NULL, linkid, modval, &modval_cnt,
3160                           media, 0, &perm_flags);
3161     if (status != DLADM_STATUS_OK)
3162         goto done;
3163
3164     for (i = 0; i < modval_cnt; i++) {
3165         if (strcasecmp(*prop_val, modval[i]) == 0) {
3166             vdp->vd_val = (uintptr_t)(uint_t)
3167                           (atof(*prop_val) * 2);
3168             status = DLADM_STATUS_OK;
3169             break;
3170         }
3171     }
3172     if (i == modval_cnt)
3173         status = DLADM_STATUS_BADVAL;
3174 done:
3175     free(buf);
3176     return (status);
3177 }
3178
3179 static dladm_status_t
3180 get_phyconf(dladm_handle_t handle, datalink_id_t linkid, void *buf,
3181              int buflen)
3182 {
3183     return (i_dladm_wlan_param(handle, linkid, buf, MAC_PROP_WL_PHY_CONFIG,
3184                                buflen, B_FALSE));
3185 }
3186
3187 /* ARGSUSED */
3188 static dladm_status_t
3189 get_channel(dladm_handle_t handle, prop_desc_t *pdp,
3190               datalink_id_t linkid, char **prop_val, uint_t *val_cntp,
3191               datalink_media_t media, uint_t flags, uint_t *perm_flags)
3192 {
3193     uint32_t channel;
3194     char buf[WLDP_BUFSIZE];
3195     dladm_status_t status;
3196     wl_phy_conf_t wl_phy_conf;
3197
3198     if ((status = get_phyconf(handle, linkid, buf, sizeof (buf)))
3199         != DLADM_STATUS_OK)
3200         return (status);
3201
3202     (void) memcpy(&wl_phy_conf, buf, sizeof (wl_phy_conf));
3203     if (!i_dladm_wlan_convert_chan(&wl_phy_conf, &channel))
3204         return (DLADM_STATUS_NOTFOUND);
3205
3206     (void) sprintf(*prop_val, DLADM_STRSIZE, "%u", channel);
3207     *val_cntp = 1;
3208     *perm_flags = MAC_PROP_PERM_READ;
3209     return (DLADM_STATUS_OK);
3210 }
3211
3212 /* ARGSUSED */
3213 static dladm_status_t
3214 get_powermode(dladm_handle_t handle, prop_desc_t *pdp,
3215                 datalink_id_t linkid, char **prop_val, uint_t *val_cntp,
3216

```

```

3216     datalink_media_t media, uint_t flags, uint_t *perm_flags)
3217 {
3218     wl_ps_mode_t mode;
3219     const char *s;
3220     char buf[WLDP_BUFSIZE];
3221     dladm_status_t status;
3222
3223     if ((status = i_dladm_wlan_param(handle, linkid, buf,
3224         MAC_PROP_WL_POWER_MODE, sizeof (buf), B_FALSE)) != DLADM_STATUS_OK)
3225         return (status);
3226
3227     (void) memcpy(&mode, buf, sizeof (mode));
3228     switch (mode.wl_ps_mode) {
3229     case WL_PM_AM:
3230         s = "off";
3231         break;
3232     case WL_PM_MPS:
3233         s = "max";
3234         break;
3235     case WL_PM_FAST:
3236         s = "fast";
3237         break;
3238     default:
3239         return (DLADM_STATUS_NOTFOUND);
3240     }
3241     (void) sprintf(*prop_val, DLADM_STRSIZE, "%s", s);
3242     *val_cnt = 1;
3243     *perm_flags = MAC_PROP_PERM_RW;
3244     return (DLADM_STATUS_OK);
3245 }
3246 /* ARGSUSED */
3247 static dladm_status_t
3248 set_powermode(dladm_handle_t handle, prop_desc_t *pdp,
3249     datalink_id_t linkid, val_desc_t *vdp, uint_t val_cnt, uint_t flags,
3250     datalink_media_t media)
3251 {
3252     dladm_wlan_powermode_t powermode = vdp->vd_val;
3253     wl_ps_mode_t ps_mode;
3254
3255     if (val_cnt != 1)
3256         return (DLADM_STATUS_BADVALCNT);
3257
3258     (void) memset(&ps_mode, 0xff, sizeof (ps_mode));
3259
3260     switch (powermode) {
3261     case DLADM_WLAN_PM_OFF:
3262         ps_mode.wl_ps_mode = WL_PM_AM;
3263         break;
3264     case DLADM_WLAN_PM_MAX:
3265         ps_mode.wl_ps_mode = WL_PM_MPS;
3266         break;
3267     case DLADM_WLAN_PM_FAST:
3268         ps_mode.wl_ps_mode = WL_PM_FAST;
3269         break;
3270     default:
3271         return (DLADM_STATUS_NOTSUP);
3272     }
3273     return (i_dladm_wlan_param(handle, linkid, &ps_mode,
3274         MAC_PROP_WL_POWER_MODE, sizeof (ps_mode), B_TRUE));
3275 }
3276 /* ARGSUSED */
3277 static dladm_status_t
3278 get_radio(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3279     char **prop_val, uint_t *val_cnt, datalink_media_t media,

```

```

3282     uint_t flags, uint_t *perm_flags)
3283 {
3284     wl_radio_t radio;
3285     const char *s;
3286     char buf[WLDP_BUFSIZE];
3287     dladm_status_t status;
3288
3289     if ((status = i_dladm_wlan_param(handle, linkid, buf,
3290         MAC_PROP_WL_RADIO, sizeof (buf), B_FALSE)) != DLADM_STATUS_OK)
3291         return (status);
3292
3293     (void) memcpy(&radio, buf, sizeof (radio));
3294     switch (radio) {
3295     case B_TRUE:
3296         s = "on";
3297         break;
3298     case B_FALSE:
3299         s = "off";
3300         break;
3301     default:
3302         return (DLADM_STATUS_NOTFOUND);
3303     }
3304     (void) sprintf(*prop_val, DLADM_STRSIZE, "%s", s);
3305     *val_cnt = 1;
3306     *perm_flags = MAC_PROP_PERM_RW;
3307     return (DLADM_STATUS_OK);
3308 }
3309 /* ARGSUSED */
3310 static dladm_status_t
3311 set_radio(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3312     val_desc_t *vdp, uint_t val_cnt, uint_t flags, datalink_media_t media)
3313 {
3314     dladm_wlan_radio_t radio = vdp->vd_val;
3315     wl_radio_t r;
3316
3317     if (val_cnt != 1)
3318         return (DLADM_STATUS_BADVALCNT);
3319
3320     switch (radio) {
3321     case DLADM_WLAN_RADIO_ON:
3322         r = B_TRUE;
3323         break;
3324     case DLADM_WLAN_RADIO_OFF:
3325         r = B_FALSE;
3326         break;
3327     default:
3328         return (DLADM_STATUS_NOTSUP);
3329     }
3330     return (i_dladm_wlan_param(handle, linkid, &r, MAC_PROP_WL_RADIO,
3331         sizeof (r), B_TRUE));
3332 }
3333 }
3334 /* ARGSUSED */
3335 static dladm_status_t
3336 check_hoplimit(dladm_handle_t handle, prop_desc_t *pdp,
3337     datalink_id_t linkid, char **prop_val, uint_t *val_cntp, uint_t flags,
3338     val_desc_t **vdp, datalink_media_t media)
3339 {
3340     int32_t hlim;
3341     char *ep;
3342     uint_t val_cntp = *val_cntp;
3343     val_desc_t *vdp = *vdp;
3344
3345     if (val_cntp != 1)
3346         return (DLADM_STATUS_BADVALCNT);

```

```

3349     errno = 0;
3350     hlim = strtol(*prop_val, &ep, 10);
3351     if (errno != 0 || ep == *prop_val || hlim < 1 ||
3352         hlim > (int32_t)UINT8_MAX)
3353         return (DLADM_STATUS_BADVAL);
3354     vdp->vd_val = hlim;
3355     return (DLADM_STATUS_OK);
3356 }

3358 /* ARGSUSED */
3359 static dladm_status_t
3360 check_encaplim(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3361                 char **prop_val, uint_t *val_cntp, uint_t flags, val_desc_t **vdpp,
3362                 datalink_media_t media)
3363 {
3364     int32_t      elim;
3365     char        *ep;
3366     uint_t       val_cnt = *val_cntp;
3367     val_desc_t   *vdp = *vdpp;

3369     if (media != DL_IPV6)
3370         return (DLADM_STATUS_BADARG);

3372     if (val_cnt != 1)
3373         return (DLADM_STATUS_BADVALCNT);

3375     errno = 0;
3376     elim = strtol(*prop_val, &ep, 10);
3377     if (errno != 0 || ep == *prop_val || elim < 0 ||
3378         elim > (int32_t)UINT8_MAX)
3379         return (DLADM_STATUS_BADVAL);
3380     vdp->vd_val = elim;
3381     return (DLADM_STATUS_OK);
3382 }

3384 static dladm_status_t
3385 i_dladm_set_linkprop_db(dladm_handle_t handle, datalink_id_t linkid,
3386                           const char *prop_name, char **prop_val, uint_t val_cnt)
3387 {
3388     char        buf[MAXLINELEN];
3389     int         i;
3390     dladm_conf_t conf;
3391     dladm_status_t status;

3393     status = dladm_open_conf(handle, linkid, &conf);
3394     if (status != DLADM_STATUS_OK)
3395         return (status);

3397     /*
3398      * reset case.
3399      */
3400     if (val_cnt == 0) {
3401         status = dladm_unset_conf_field(handle, conf, prop_name);
3402         if (status == DLADM_STATUS_OK)
3403             status = dladm_write_conf(handle, conf);
3404         goto done;
3405     }

3407     buf[0] = '\0';
3408     for (i = 0; i < val_cnt; i++) {
3409         (void) strlcat(buf, prop_val[i], MAXLINELEN);
3410         if (i != val_cnt - 1)
3411             (void) strlcat(buf, ",", MAXLINELEN);
3412     }

```

```

3414     status = dladm_set_conf_field(handle, conf, prop_name, DLADM_TYPE_STR,
3415                                   buf);
3416     if (status == DLADM_STATUS_OK)
3417         status = dladm_write_conf(handle, conf);

3419 done:
3420     dladm_destroy_conf(handle, conf);
3421     return (status);
3422 }

3424 static dladm_status_t
3425 i_dladm_get_linkprop_db(dladm_handle_t handle, datalink_id_t linkid,
3426                           const char *prop_name, char **prop_val, uint_t *val_cntp)
3427 {
3428     char        buf[MAXLINELEN], *str;
3429     uint_t      cnt = 0;
3430     dladm_conf_t conf;
3431     dladm_status_t status;

3433     status = dladm_getsnap_conf(handle, linkid, &conf);
3434     if (status != DLADM_STATUS_OK)
3435         return (status);

3437     status = dladm_get_conf_field(handle, conf, prop_name, buf, MAXLINELEN);
3438     if (status != DLADM_STATUS_OK)
3439         goto done;

3441     str = strtok(buf, ",");
3442     while (str != NULL) {
3443         if (cnt == *val_cntp) {
3444             status = DLADM_STATUS_TOOSMALL;
3445             goto done;
3446         }
3447         (void) strlcpy(prop_val[cnt++], str, DLADM_PROP_VAL_MAX);
3448         str = strtok(NULL, ",");
3449     }

3451     *val_cntp = cnt;

3453 done:
3454     dladm_destroy_conf(handle, conf);
3455     return (status);
3456 }

3458 /*
3459  * Walk persistent private link properties of a link.
3460  */
3461 static dladm_status_t
3462 i_dladm_walk_linkprop_priv_db(dladm_handle_t handle, datalink_id_t linkid,
3463                                 void *arg, int (*func)(dladm_handle_t, datalink_id_t, const char *, void *))
3464 {
3465     dladm_status_t    status;
3466     dladm_conf_t      conf;
3467     char            last_attr[MAXLINKATTRLEN];
3468     char            attr[MAXLINKATTRLEN];
3469     char            attrval[MAXLINKATTRVALLEN];
3470     size_t          attrsz;

3472     if (linkid == DATALINK_INVALID_LINKID || func == NULL)
3473         return (DLADM_STATUS_BADARG);

3475     status = dladm_getsnap_conf(handle, linkid, &conf);
3476     if (status != DLADM_STATUS_OK)
3477         return (status);

3479     last_attr[0] = '\0';

```

```

3480     while ((status = dladm_getnext_conf_linkprop(handle, conf, last_attr,
3481           attr, attrval, MAXLINKATTRVALLEN, &attrsz)) == DLADM_STATUS_OK) {
3482         if (attr[0] == '_') {
3483             if (func(handle, linkid, attr, arg) ==
3484                 DLADM_WALK_TERMINATE)
3485                 break;
3486         }
3487         (void) strlcpy(last_attr, attr, MAXLINKATTRLEN);
3488     }
3489
3490     dladm_destroy_conf(handle, conf);
3491     return (DLADM_STATUS_OK);
3492 }
3493
3494 static link_attr_t *
3495 dladm_name2prop(const char *prop_name)
3496 {
3497     link_attr_t *p;
3498
3499     for (p = link_attr; p->pp_id != MAC_PROP_PRIVATE; p++) {
3500         if (strcmp(p->pp_name, prop_name) == 0)
3501             break;
3502     }
3503     return (p);
3504 }
3505
3506 static link_attr_t *
3507 dladm_id2prop(mac_prop_id_t propid)
3508 {
3509     link_attr_t *p;
3510
3511     for (p = link_attr; p->pp_id != MAC_PROP_PRIVATE; p++) {
3512         if (p->pp_id == propid)
3513             break;
3514     }
3515     return (p);
3516 }
3517
3518 static dld_ioc_macprop_t *
3519 i_dladm_buf_allocImpl(size_t valsiz, datalink_id_t linkid,
3520   const char *prop_name, mac_prop_id_t propid, uint_t flags,
3521   dladm_status_t *status)
3522 {
3523     int dsiz;
3524     dld_ioc_macprop_t *dip;
3525
3526     *status = DLADM_STATUS_OK;
3527     dsiz = MAC_PROP_BUFSIZE(valsiz);
3528     dip = malloc(dsiz);
3529     if (dip == NULL) {
3530         *status = DLADM_STATUS_NOMEM;
3531         return (NULL);
3532     }
3533     bzero(dip, dsiz);
3534     dip->pr_valsiz = valsiz;
3535     (void) strlcpy(dip->pr_name, prop_name, sizeof (dip->pr_name));
3536     dip->pr_linkid = linkid;
3537     dip->pr_num = propid;
3538     dip->pr_flags = flags;
3539     return (dip);
3540 }
3541
3542 static dld_ioc_macprop_t *
3543 i_dladm_buf_allocByName(size_t valsiz, datalink_id_t linkid,
3544   const char *prop_name, uint_t flags, dladm_status_t *status)
3545 {

```

```

3546     link_attr_t *p;
3547
3548     p = dladm_name2prop(prop_name);
3549     valsiz = MAX(p->pp_valsiz, valsiz);
3550     return (i_dladm_buf_allocImpl(valsiz, linkid, prop_name, p->pp_id,
3551       flags, status));
3552 }
3553
3554 static dld_ioc_macprop_t *
3555 i_dladm_buf_allocBy_id(size_t valsiz, datalink_id_t linkid,
3556   mac_prop_id_t propid, uint_t flags, dladm_status_t *status)
3557 {
3558     link_attr_t *p;
3559
3560     p = dladm_id2prop(propid);
3561     valsiz = MAX(p->pp_valsiz, valsiz);
3562     return (i_dladm_buf_allocImpl(valsiz, linkid, p->pp_name, propid,
3563       flags, status));
3564 }
3565
3566 /* ARGSUSED */
3567 static dladm_status_t
3568 set_public_prop(dladm_handle_t handle, prop_desc_t *pdp,
3569   datalink_id_t linkid, val_desc_t *vdp, uint_t val_cnt, uint_t flags,
3570   datalink_media_t media)
3571 {
3572     dld_ioc_macprop_t      *dip;
3573     dladm_status_t          status = DLADM_STATUS_OK;
3574     uint8_t                 u8;
3575     uint16_t                u16;
3576     uint32_t                u32;
3577     void                    *val;
3578
3579     dip = i_dladm_buf_allocByName(0, linkid, pdp->pd_name, 0, &status);
3580     if (dip == NULL)
3581         return (status);
3582
3583     if (pdp->pd_flags & PD_CHECK_ALLOC)
3584         val = (void *)vdp->vd_val;
3585     else {
3586         /*
3587          * Currently all 1/2/4-byte size properties are byte/word/int.
3588          * No need (yet) to distinguish these from arrays of same size.
3589          */
3590         switch (dip->pr_valsiz) {
3591             case 1:
3592                 u8 = vdp->vd_val;
3593                 val = &u8;
3594                 break;
3595             case 2:
3596                 u16 = vdp->vd_val;
3597                 val = &u16;
3598                 break;
3599             case 4:
3600                 u32 = vdp->vd_val;
3601                 val = &u32;
3602                 break;
3603             default:
3604                 val = &vdp->vd_val;
3605                 break;
3606         }
3607     }
3608
3609     if (val != NULL)
3610         (void) memcpy(dip->pr_val, val, dip->pr_valsiz);
3611     else

```

```

3612         dip->pr_valsize = 0;
3614     status = i_dladm_macprop(handle, dip, B_TRUE);
3616 done:
3617     free(dip);
3618     return (status);
3619 }

3621 dladm_status_t
3622 i_dladm_macprop(dladm_handle_t handle, void *dip, boolean_t set)
3623 {
3624     dladm_status_t status = DLADM_STATUS_OK;
3626     if (ioctl(dladm_dld_fd(handle),
3627             (set ? DLDIOC_SETMACPROP : DLDIOC_GETMACPROP), dip))
3628         status = dladm_errno2status(errno);
3630     return (status);
3631 }

3633 static dladm_status_t
3634 i_dladm_get_public_prop(dladm_handle_t handle, datalink_id_t linkid,
3635     char *prop_name, uint_t flags, uint_t *perm_flags, void *arg, size_t size)
3636 {
3637     dld_ioc_macprop_t      *dip;
3638     dladm_status_t          status;
3640     dip = i_dladm_buf_alloc_by_name(0, linkid, prop_name, flags, &status);
3641     if (dip == NULL)
3642         return (DLADM_STATUS_NOMEM);
3644     status = i_dladm_macprop(handle, dip, B_FALSE);
3645     if (status != DLADM_STATUS_OK) {
3646         free(dip);
3647         return (status);
3648     }
3650     if (perm_flags != NULL)
3651         *perm_flags = dip->pr_perm_flags;
3653     if (arg != NULL)
3654         (void) memcpy(arg, dip->pr_val, size);
3655     free(dip);
3656     return (DLADM_STATUS_OK);
3657 }

3659 /* ARGSUSED */
3660 static dladm_status_t
3661 check_uint32(dladm_handle_t handle, prop_desc_t *pdp,
3662     datalink_id_t linkid, char **prop_val, uint_t *val_cntp, uint_t flags,
3663     val_desc_t **vp, datalink_media_t media)
3664 {
3665     uint_t          val_cnt = *val_cntp;
3666     val_desc_t      *v = *vp;
3668     if (val_cnt != 1)
3669         return (DLADM_STATUS_BADVAL);
3670     v->vd_val = strtoul(prop_val[0], NULL, 0);
3671     return (DLADM_STATUS_OK);
3672 }

3674 /* ARGSUSED */
3675 static dladm_status_t
3676 get_duplex(dladm_handle_t handle, prop_desc_t *pdp,
3677     datalink_id_t linkid, char **prop_val, uint_t *val_cnt,

```

```

3678     datalink_media_t media, uint_t flags, uint_t *perm_flags)
3679 {
3680     link_duplex_t    link_duplex;
3681     dladm_status_t   status;
3683     if ((status = dladm_get_single_mac_stat(handle, linkid, "link_duplex",
3684         KSTAT_DATA_UINT32, &link_duplex)) != 0)
3685         return (status);
3687     switch (link_duplex) {
3688     case LINK_DUPLEX_FULL:
3689         (void) strcpy(*prop_val, "full");
3690         break;
3691     case LINK_DUPLEX_HALF:
3692         (void) strcpy(*prop_val, "half");
3693         break;
3694     default:
3695         (void) strcpy(*prop_val, "unknown");
3696         break;
3697     }
3698     *val_cntp = 1;
3699     return (DLADM_STATUS_OK);
3700 }

3702 /* ARGSUSED */
3703 static dladm_status_t
3704 get_speed(dladm_handle_t handle, prop_desc_t *pdp, datalink_id_t linkid,
3705     char **prop_val, uint_t *val_cntp, datalink_media_t media, uint_t flags,
3706     uint_t *perm_flags)
3707 {
3708     uint64_t          ifspeed = 0;
3709     dladm_status_t   status;
3711     if ((status = dladm_get_single_mac_stat(handle, linkid, "ifspeed",
3712         KSTAT_DATA_UINT64, &ifspeed)) != 0)
3713         return (status);
3715     if ((ifspeed % 1000000) != 0) {
3716         (void) sprintf(*prop_val, DLADM_PROP_VAL_MAX,
3717             "%lld", ifspeed / (float)1000000); /* Mbps */
3718     } else {
3719         (void) sprintf(*prop_val, DLADM_PROP_VAL_MAX,
3720             "%llu", ifspeed / 1000000); /* Mbps */
3721     }
3722     *val_cntp = 1;
3723     *perm_flags = MAC_PROP_PERM_READ;
3724     return (DLADM_STATUS_OK);
3725 }

3727 /* ARGSUSED */
3728 static dladm_status_t
3729 get_link_state(dladm_handle_t handle, prop_desc_t *pdp,
3730     datalink_id_t linkid, char **prop_val, uint_t *val_cntp,
3731     datalink_media_t media, uint_t flags, uint_t *perm_flags)
3732 {
3733     link_state_t      link_state;
3734     dladm_status_t   status;
3736     status = dladm_get_state(handle, linkid, &link_state);
3737     if (status != DLADM_STATUS_OK)
3738         return (status);
3740     switch (link_state) {
3741     case LINK_STATE_UP:
3742         (void) strcpy(*prop_val, "up");
3743         break;

```

```

3744     case LINK_STATE_DOWN:
3745         (void) strcpy(*prop_val, "down");
3746         break;
3747     default:
3748         (void) strcpy(*prop_val, "unknown");
3749         break;
3750     }
3751     *val_cnt = 1;
3752     *perm_flags = MAC_PROP_PERM_READ;
3753     return (DLADM_STATUS_OK);
3754 }

3756 /* ARGSUSED */
3757 static dladm_status_t
3758 get_binary(dladm_handle_t handle, prop_desc_t *pdp,
3759             datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
3760             datalink_media_t media, uint_t flags, uint_t *perm_flags)
3761 {
3762     dladm_status_t status;
3763     uint_t v = 0;

3765     status = i_dladm_get_public_prop(handle, linkid, pdp->pd_name, flags,
3766                                     perm_flags, &v, sizeof(v));
3767     if (status != DLADM_STATUS_OK)
3768         return (status);

3770     (void) sprintf(*prop_val, DLADM_PROP_VAL_MAX, "%d", (uint_t)(v > 0));
3771     *val_cnt = 1;
3772     return (DLADM_STATUS_OK);
3773 }

3775 /* ARGSUSED */
3776 static dladm_status_t
3777 get_uint32(dladm_handle_t handle, prop_desc_t *pdp,
3778             datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
3779             datalink_media_t media, uint_t flags, uint_t *perm_flags)
3780 {
3781     dladm_status_t status;
3782     uint32_t v = 0;

3784     status = i_dladm_get_public_prop(handle, linkid, pdp->pd_name, flags,
3785                                     perm_flags, &v, sizeof(v));
3786     if (status != DLADM_STATUS_OK)
3787         return (status);

3789     (void) sprintf(*prop_val, DLADM_PROP_VAL_MAX, "%ld", v);
3790     *val_cnt = 1;
3791     return (DLADM_STATUS_OK);
3792 }

3794 /* ARGSUSED */
3795 static dladm_status_t
3796 get_range(dladm_handle_t handle, prop_desc_t *pdp,
3797             datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
3798             datalink_media_t media, uint_t flags, uint_t *perm_flags)
3799 {
3800     dld_ioc_macprop_t *dip;
3801     dladm_status_t status = DLADM_STATUS_OK;
3802     size_t sz;
3803     uint_t rcount;
3804     mac_propval_range_t *rangep;

3806     /*
3807      * As caller we don't know number of value ranges, the driver
3808      * supports. To begin with we assume that number to be 1. If the
3809      * buffer size is insufficient, driver returns back with the

```

```

3810         * actual count of value ranges. See mac.h for more details.
3811         */
3812         sz = sizeof (mac_propval_range_t);
3813         rcount = 1;
3814     retry:
3815         if ((dip = i_dladm_buf_alloc_by_name(sz, linkid, pdp->pd_name, flags,
3816                                             &status)) == NULL)
3817             return (status);

3819     rangep = (mac_propval_range_t *) (void *) &dip->pr_val;
3820     rangep->mpr_count = rcount;

3822     status = i_dladm_macprop(handle, dip, B_FALSE);
3823     if (status != DLADM_STATUS_OK) {
3824         if (status == DLADM_STATUS_TOOSMALL) {
3825             int err;

3827             if ((err = i_dladm_range_size(rangep, &sz, &rcount))
3828                 == 0) {
3829                 free(dip);
3830                 goto retry;
3831             } else {
3832                 status = dladm_errno2status(err);
3833             }
3834             free(dip);
3835             return (status);
3836         }
3837     }

3839     if (rangep->mpr_count == 0) {
3840         *val_cnt = 1;
3841         (void) sprintf(prop_val[0], DLADM_PROP_VAL_MAX, "--");
3842         goto done;
3843     }

3845     switch (rangep->mpr_type) {
3846     case MAC_PROPVAL_UINT32: {
3847         mac_propval_uint32_range_t *ur;
3848         uint_t count = rangep->mpr_count, i;

3850         ur = &rangep->mpr_range_uint32[0];
3851         for (i = 0; i < count; i++, ur++) {
3852             if (ur->mpur_min == ur->mpur_max) {
3853                 (void) sprintf(prop_val[i], DLADM_PROP_VAL_MAX,
3854                               "%ld", ur->mpur_min);
3855             } else {
3856                 (void) sprintf(prop_val[i], DLADM_PROP_VAL_MAX,
3857                               "%ld-%ld", ur->mpur_min, ur->mpur_max);
3858             }
3859         }
3860         *val_cnt = count;
3861         break;
3862     }
3863     default:
3864         status = DLADM_STATUS_BADARG;
3865         break;
3866     }
3867 done:
3868     free(dip);
3869     return (status);
3870 }

3873 /* ARGSUSED */
3874 static dladm_status_t
3875 get_tagmode(dladm_handle_t handle, prop_desc_t *pdp,

```

```

3876     datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
3877     datalink_media_t media, uint_t flags, uint_t *perm_flags)
3878 {
3879     link_tagmode_t mode;
3880     dladm_status_t status;
3881
3882     status = i_dladm_get_public_prop(handle, linkid, pdp->pd_name, flags,
3883         perm_flags, &mode, sizeof(mode));
3884     if (status != DLADM_STATUS_OK)
3885         return (status);
3886
3887     switch (mode) {
3888     case LINK_TAGMODE_NORMAL:
3889         (void) strlcpy(*prop_val, "normal", DLADM_PROP_VAL_MAX);
3890         break;
3891     case LINK_TAGMODE_VLANONLY:
3892         (void) strlcpy(*prop_val, "vlanonly", DLADM_PROP_VAL_MAX);
3893         break;
3894     default:
3895         (void) strlcpy(*prop_val, "unknown", DLADM_PROP_VAL_MAX);
3896     }
3897     *val_cnt = 1;
3898     return (DLADM_STATUS_OK);
3899 }
3900 /* ARGSUSED */
3901 static dladm_status_t
3902 get_flowctrl(dladm_handle_t handle, prop_desc_t *pdp,
3903     datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
3904     datalink_media_t media, uint_t flags, uint_t *perm_flags)
3905 {
3906     link_flowctrl_t v;
3907     dladm_status_t status;
3908
3909     status = i_dladm_get_public_prop(handle, linkid, pdp->pd_name, flags,
3910         perm_flags, &v, sizeof(v));
3911     if (status != DLADM_STATUS_OK)
3912         return (status);
3913
3914     switch (v) {
3915     case LINK_FLOWCTRL_NONE:
3916         (void) sprintf(*prop_val, "no");
3917         break;
3918     case LINK_FLOWCTRL_RX:
3919         (void) sprintf(*prop_val, "rx");
3920         break;
3921     case LINK_FLOWCTRL_TX:
3922         (void) sprintf(*prop_val, "tx");
3923         break;
3924     case LINK_FLOWCTRL_BI:
3925         (void) sprintf(*prop_val, "bi");
3926         break;
3927     }
3928     *val_cnt = 1;
3929     return (DLADM_STATUS_OK);
3930 }
3931 }
3932 /* ARGSUSED */
3933 static dladm_status_t
3934 i_dladm_set_private_prop(dladm_handle_t handle, datalink_id_t linkid,
3935     const char *prop_name, char **prop_val, uint_t val_cnt, uint_t flags)
3936 {
3937     int i, selen;
3938     bufsize = 0;

```

```

3942     dld_ioc_macprop_t *dip = NULL;
3943     uchar_t *dp;
3944     link_attr_t *p;
3945     dladm_status_t status = DLADM_STATUS_OK;
3946
3947     if ((prop_name == NULL && prop_val != NULL) ||
3948         (prop_val != NULL && val_cnt == 0))
3949         return (DLADM_STATUS_BADARG);
3950     p = dladm_name2prop(prop_name);
3951     if (p->pp_id != MAC_PROP_PRIVATE)
3952         return (DLADM_STATUS_BADARG);
3953
3954     if (!(flags & DLADM_OPT_ACTIVE))
3955         return (DLADM_STATUS_OK);
3956
3957     /*
3958      * private properties: all parsing is done in the kernel.
3959      * allocate a enough space for each property + its separator (',').
3960      */
3961     for (i = 0; i < val_cnt; i++) {
3962         bufsize += strlen(prop_val[i]) + 1;
3963     }
3964
3965     if (prop_val == NULL) {
3966         /*
3967          * getting default value. so use more buffer space.
3968          */
3969         bufsize += DLADM_PROP_BUF_CHUNK;
3970     }
3971
3972     dip = i_dladm_buf_alloc_by_name(bufsize + 1, linkid, prop_name,
3973         (prop_val != NULL ? 0 : DLD_PROP_DEFAULT), &status);
3974     if (dip == NULL)
3975         return (status);
3976
3977     dp = (uchar_t *)dip->pr_val;
3978     selen = 0;
3979
3980     if (prop_val == NULL) {
3981         status = i_dladm_macprop(handle, dip, B_FALSE);
3982         dip->pr_flags = 0;
3983     } else {
3984         for (i = 0; i < val_cnt; i++) {
3985             int plen = 0;
3986
3987             plen = strlen(prop_val[i]);
3988             bcopy(prop_val[i], dp, plen);
3989             selen += plen;
3990
3991             /*
3992              * add a "," separator and update dp.
3993              */
3994             if (i != (val_cnt - 1))
3995                 dp[slen++] = ',';
3996             dp += (plen + 1);
3997         }
3998     }
3999     if (status == DLADM_STATUS_OK)
4000         status = i_dladm_macprop(handle, dip, B_TRUE);
4001
4002     free(dip);
4003     return (status);
4004
4005     static dladm_status_t
4006     i_dladm_get_priv_prop(dladm_handle_t handle, datalink_id_t linkid,
4007         const char *prop_name, char **prop_val, uint_t *val_cnt,

```

```

4008     dladm_prop_type_t type, uint_t dld_flags)
4009 {
4010     dladm_status_t status = DLADM_STATUS_OK;
4011     dld_ioc_macprop_t *dip = NULL;
4012     link_attr_t *p;
4013
4014     if ((prop_name == NULL && prop_val != NULL) ||
4015         (prop_val != NULL && val_cnt == 0))
4016         return (DLADM_STATUS_BADARG);
4017
4018     p = dladm_name2prop(prop_name);
4019     if (p->pp_id != MAC_PROP_PRIVATE)
4020         return (DLADM_STATUS_BADARG);
4021
4022     /*
4023      * private properties: all parsing is done in the kernel.
4024      */
4025     dip = i_dladm_buf_alloc_by_name(DLADM_PROP_BUF_CHUNK, linkid, prop_name,
4026                                     dld_flags, &status);
4027     if (dip == NULL)
4028         return (status);
4029
4030     if ((status = i_dladm_macprop(handle, dip, B_FALSE)) ==
4031         DLADM_STATUS_OK) {
4032         if (type == DLADM_PROP_VAL_PERM) {
4033             (void) dladm_perm2str(dip->pr_perm_flags, *prop_val);
4034         } else if (type == DLADM_PROP_VAL_MODIFIABLE) {
4035             *prop_val[0] = '\0';
4036         } else {
4037             (void) strncpy(*prop_val, dip->pr_val,
4038                           DLADM_PROP_VAL_MAX);
4039         }
4040         *val_cnt = 1;
4041     } else if ((status == DLADM_STATUS_NOTSUP) &&
4042                (type == DLADM_PROP_VAL_CURRENT)) {
4043         status = DLADM_STATUS_NOTFOUND;
4044     }
4045     free(dip);
4046     return (status);
4047 }
4048
4049 static dladm_status_t
4050 i_dladm_getset_defval(dladm_handle_t handle, prop_desc_t *pdp,
4051                        datalink_id_t linkid, datalink_media_t media, uint_t flags)
4052 {
4053     dladm_status_t status;
4054     char **prop_vals = NULL, *buf;
4055     size_t bufsize;
4056     uint_t cnt;
4057     int i;
4058     uint_t perm_flags;
4059
4060     /*
4061      * Allocate buffer needed for prop_vals array. We can have at most
4062      * DLADM_MAX_PROP_VALCNT char *prop_vals[] entries, where
4063      * each entry has max size DLADM_PROP_VAL_MAX
4064      */
4065     bufsize =
4066         (sizeof (char *) + DLADM_PROP_VAL_MAX) * DLADM_MAX_PROP_VALCNT;
4067     buf = malloc(bufsize);
4068     prop_vals = (char **)(void *)buf;
4069     for (i = 0; i < DLADM_MAX_PROP_VALCNT; i++) {
4070         prop_vals[i] = buf +
4071             sizeof (char *) * DLADM_MAX_PROP_VALCNT +
4072             i * DLADM_PROP_VAL_MAX;

```

```

4074     }
4075
4076     /*
4077      * For properties which have pdp->pd_defval.vd_name as a non-empty
4078      * string, the "" itself is used to reset the property (exceptions
4079      * are zone and autopush, which populate vdp->vd_val). So
4080      * libdladm can copy pdp->pd_defval over to the val_desc_t passed
4081      * down on the setprop using the global values in the table. For
4082      * other cases (vd_name is ""), doing reset-linkprop will cause
4083      * libdladm to do a getprop to find the default value and then do
4084      * a setprop to reset the value to default.
4085     */
4086     status = pdp->pd_get(handle, pdp, linkid, prop_vals, &cnt, media,
4087                           DLD_PROP_DEFAULT, &perm_flags);
4088     if (status == DLADM_STATUS_OK) {
4089         if (perm_flags == MAC_PROP_PERM_RW) {
4090             status = i_dladm_set_single_prop(handle, linkid,
4091                                             pdp->pd_class, media, pdp, prop_vals, cnt, flags);
4092         }
4093         else
4094             status = DLADM_STATUS_NOTSUP;
4095     }
4096     free(buf);
4097     return (status);
4098 }
4099
4100 /* ARGSUSED */
4101 static dladm_status_t
4102 get_stp(dladm_handle_t handle, struct prop_desc *pd, datalink_id_t linkid,
4103          char **prop_val, uint_t *val_cnt, datalink_media_t media, uint_t flags,
4104          uint_t *perm_flags)
4105 {
4106     const bridge_public_prop_t *bpp;
4107     dladm_status_t retv;
4108     int val, i;
4109
4110     if (flags != 0)
4111         return (DLADM_STATUS_NOTSUP);
4112     *perm_flags = MAC_PROP_PERM_RW;
4113     *val_cnt = 1;
4114     for (bpp = bridge_prop; bpp->bpp_name != NULL; bpp++)
4115         if (strcmp(bpp->bpp_name, pd->pd_name) == 0)
4116             break;
4117     retv = dladm_bridge_get_port_cfg(handle, linkid, bpp->bpp_code, &val);
4118     /* If the daemon isn't running, then return the persistent value */
4119     if (retv == DLADM_STATUS_NOTFOUND) {
4120         if (i_dladm_get_linkprop_db(handle, linkid, pd->pd_name,
4121                                     prop_val, val_cnt) != DLADM_STATUS_OK)
4122             (void) strlcpy(*prop_val, pd->pd_defval.vd_name,
4123                           DLADM_PROP_VAL_MAX);
4124         return (DLADM_STATUS_OK);
4125     }
4126     if (retv != DLADM_STATUS_OK) {
4127         (void) strlcpy(*prop_val, "?", DLADM_PROP_VAL_MAX);
4128         return (retv);
4129     }
4130     if (val == pd->pd_defval.vd_val && pd->pd_defval.vd_name[0] != '\0') {
4131         (void) strlcpy(*prop_val, pd->pd_defval.vd_name,
4132                       DLADM_PROP_VAL_MAX);
4133         return (DLADM_STATUS_OK);
4134     }
4135     for (i = 0; i < pd->pd_noptval; i++) {
4136         if (val == pd->pd_optval[i].vd_val) {
4137             (void) strlcpy(*prop_val, pd->pd_optval[i].vd_name,
4138                           DLADM_PROP_VAL_MAX);
4139             return (DLADM_STATUS_OK);
4140         }
4141     }

```

```

4140         }
4141     }
4142     (void) snprintf(*prop_val, DLADM_PROP_VAL_MAX, "%u", (unsigned)val);
4143     return (DLADM_STATUS_OK);
4144 }

4146 /* ARGSUSED1 */
4147 static dladm_status_t
4148 set_stp_prop(dladm_handle_t handle, prop_desc_t *pd, datalink_id_t linkid,
4149   val_desc_t *vdp, uint_t val_cnt, uint_t flags, datalink_media_t media)
4150 {
4151     /*
4152      * Special case for mcheck: the daemon resets the value to zero, and we
4153      * don't want the daemon to refresh itself; it leads to deadlock.
4154      */
4155     if (flags & DLADM_OPT_NOREFRESH)
4156         return (DLADM_STATUS_OK);

4158     /* Tell the running daemon, if any */
4159     return (dladm_bridge_refresh(handle, linkid));
4160 }

4162 /*
4163  * This is used only for stp_priority, stp_cost, and stp_mcheck.
4164 */
4165 /* ARGSUSED */
4166 static dladm_status_t
4167 check_stp_prop(dladm_handle_t handle, struct prop_desc *pd,
4168   datalink_id_t linkid, char **prop_val, uint_t *val_cntp, uint_t flags,
4169   val_desc_t **vdp, datalink_media_t media)
4170 {
4171     char          *cp;
4172     boolean_t      iscost;
4173     uint_t         val_cnt = *val_cntp;
4174     val_desc_t    *vdp = *vdp;

4176     if (val_cnt != 1)
4177         return (DLADM_STATUS_BADVALCNT);

4179     if (prop_val == NULL) {
4180         vdp->vd_val = 0;
4181     } else {
4182         /* Only stp_priority and stp_cost use this function */
4183         iscost = strcmp(pd->pd_name, "stp_cost") == 0;

4185         if (iscost && strcmp(prop_val[0], "auto") == 0) {
4186             /* Illegal value 0 is allowed to mean "automatic" */
4187             vdp->vd_val = 0;
4188         } else {
4189             errno = 0;
4190             vdp->vd_val = strtoul(prop_val[0], &cp, 0);
4191             if (errno != 0 || *cp != '\0')
4192                 return (DLADM_STATUS_BADVAL);
4193         }
4194     }

4196     if (iscost) {
4197         return (vdp->vd_val > 65535 ? DLADM_STATUS_BADVAL :
4198             DLADM_STATUS_OK);
4199     } else {
4200         if (vdp->vd_val > 255)
4201             return (DLADM_STATUS_BADVAL);
4202         /*
4203          * If the user is setting stp_mcheck non-zero, then (per the
4204          * IEEE management standards and UNH testing) we need to check
4205          * whether this link is part of a bridge that is running RSTP.
4206
4207
4208
4209
4210
4211
4212
4213
4214
4215
4216
4217
4218
4219
4220
4221
4222
4223
4224
4225
4226
4227
4228
4229
4230
4231
4232
4233
4234
4235
4236
4237
4238
4239
4240
4241
4242
4243
4244
4245
4246
4247
4248
4249
4250
4251
4252
4253
4254
4255
4256
4257
4258
4259
4260
4261
4262
4263
4264
4265
4266
4267
4268
4269
4270
4271
4272
4273
4274
4275
4276
4277
4278
4279
4280
4281
4282
4283
4284
4285
4286
4287
4288
4289
4290
4291
4292
4293
4294
4295
4296
4297
4298
4299
4300
4301
4302
4303
4304
4305
4306
4307
4308
4309
4310
4311
4312
4313
4314
4315
4316
4317
4318
4319
4320
4321
4322
4323
4324
4325
4326
4327
4328
4329
4330
4331
4332
4333
4334
4335
4336
4337
4338
4339
4340
4341
4342
4343
4344
4345
4346
4347
4348
4349
4350
4351
4352
4353
4354
4355
4356
4357
4358
4359
4360
4361
4362
4363
4364
4365
4366
4367
4368
4369
4370
4371
4372
4373
4374
4375
4376
4377
4378
4379
4380
4381
4382
4383
4384
4385
4386
4387
4388
4389
4390
4391
4392
4393
4394
4395
4396
4397
4398
4399
4400
4401
4402
4403
4404
4405
4406
4407
4408
4409
4410
4411
4412
4413
4414
4415
4416
4417
4418
4419
4420
4421
4422
4423
4424
4425
4426
4427
4428
4429
4430
4431
4432
4433
4434
4435
4436
4437
4438
4439
4440
4441
4442
4443
4444
4445
4446
4447
4448
4449
4450
4451
4452
4453
4454
4455
4456
4457
4458
4459
4460
4461
4462
4463
4464
4465
4466
4467
4468
4469
4470
4471
4472
4473
4474
4475
4476
4477
4478
4479
4480
4481
4482
4483
4484
4485
4486
4487
4488
4489
4490
4491
4492
4493
4494
4495
4496
4497
4498
4499
4500
4501
4502
4503
4504
4505
4506
4507
4508
4509
4510
4511
4512
4513
4514
4515
4516
4517
4518
4519
4520
4521
4522
4523
4524
4525
4526
4527
4528
4529
4530
4531
4532
4533
4534
4535
4536
4537
4538
4539
4540
4541
4542
4543
4544
4545
4546
4547
4548
4549
4550
4551
4552
4553
4554
4555
4556
4557
4558
4559
4560
4561
4562
4563
4564
4565
4566
4567
4568
4569
4570
4571
4572
4573
4574
4575
4576
4577
4578
4579
4580
4581
4582
4583
4584
4585
4586
4587
4588
4589
4590
4591
4592
4593
4594
4595
4596
4597
4598
4599
4599
4600
4601
4602
4603
4604
4605
4606
4607
4608
4609
4610
4611
4612
4613
4614
4615
4616
4617
4618
4619
4620
4621
4622
4623
4624
4625
4626
4627
4628
4629
4630
4631
4632
4633
4634
4635
4636
4637
4638
4639
4640
4641
4642
4643
4644
4645
4646
4647
4648
4649
4650
4651
4652
4653
4654
4655
4656
4657
4658
4659
4660
4661
4662
4663
4664
4665
4666
4667
4668
4669
4670
4671
4672
4673
4674
4675
4676
4677
4678
4679
4680
4681
4682
4683
4684
4685
4686
4687
4688
4689
4690
4691
4692
4693
4694
4695
4696
4697
4698
4699
4699
4700
4701
4702
4703
4704
4705
4706
4707
4708
4709
4709
4710
4711
4712
4713
4714
4715
4716
4717
4718
4719
4719
4720
4721
4722
4723
4724
4725
4726
4727
4728
4729
4729
4730
4731
4732
4733
4734
4735
4736
4737
4738
4739
4739
4740
4741
4742
4743
4744
4745
4746
4747
4748
4749
4749
4750
4751
4752
4753
4754
4755
4756
4757
4758
4759
4759
4760
4761
4762
4763
4764
4765
4766
4767
4768
4769
4769
4770
4771
4772
4773
4774
4775
4776
4777
4778
4779
4779
4780
4781
4782
4783
4784
4785
4786
4787
4788
4789
4789
4790
4791
4792
4793
4794
4795
4796
4797
4798
4799
4799
4800
4801
4802
4803
4804
4805
4806
4807
4808
4809
4809
4810
4811
4812
4813
4814
4815
4816
4817
4818
4819
4819
4820
4821
4822
4823
4824
4825
4826
4827
4828
4829
4829
4830
4831
4832
4833
4834
4835
4836
4837
4838
4839
4839
4840
4841
4842
4843
4844
4845
4846
4847
4848
4849
4849
4850
4851
4852
4853
4854
4855
4856
4857
4858
4859
4859
4860
4861
4862
4863
4864
4865
4866
4867
4868
4869
4869
4870
4871
4872
4873
4874
4875
4876
4877
4878
4879
4879
4880
4881
4882
4883
4884
4885
4886
4887
4888
4889
4889
4890
4891
4892
4893
4894
4895
4896
4897
4898
4899
4899
4900
4901
4902
4903
4904
4905
4906
4907
4908
4909
4909
4910
4911
4912
4913
4914
4915
4916
4917
4918
4919
4919
4920
4921
4922
4923
4924
4925
4926
4927
4928
4929
4929
4930
4931
4932
4933
4934
4935
4936
4937
4938
4939
4939
4940
4941
4942
4943
4944
4945
4946
4947
4948
4949
4949
4950
4951
4952
4953
4954
4955
4956
4957
4958
4959
4959
4960
4961
4962
4963
4964
4965
4966
4967
4968
4969
4969
4970
4971
4972
4973
4974
4975
4976
4977
4978
4979
4979
4980
4981
4982
4983
4984
4985
4986
4987
4988
4989
4989
4989
4990
4991
4992
4993
4994
4995
4996
4997
4998
4999
4999
5000
5001
5002
5003
5004
5005
5006
5007
5008
5009
5009
5010
5011
5012
5013
5014
5015
5016
5017
5018
5019
5019
5020
5021
5022
5023
5024
5025
5026
5027
5028
5029
5029
5030
5031
5032
5033
5034
5035
5036
5037
5038
5039
5039
5040
5041
5042
5043
5044
5045
5046
5047
5048
5049
5049
5050
5051
5052
5053
5054
5055
5056
5057
5058
5059
5059
5060
5061
5062
5063
5064
5065
5066
5067
5068
5069
5069
5070
5071
5072
5073
5074
5075
5076
5077
5078
5079
5079
5080
5081
5082
5083
5084
5085
5086
5087
5088
5089
5089
5089
5090
5091
5092
5093
5094
5095
5096
5097
5098
5099
5099
5099
5100
5101
5102
5103
5104
5105
5106
5107
5108
5109
5109
5110
5111
5112
5113
5114
5115
5116
5117
5118
5119
5119
5120
5121
5122
5123
5124
5125
5126
5127
5128
5129
5129
5130
5131
5132
5133
5134
5135
5136
5137
5138
5139
5139
5140
5141
5142
5143
5144
5145
5146
5147
5148
5149
5149
5150
5151
5152
5153
5154
5155
5156
5157
5158
5159
5159
5160
5161
5162
5163
5164
5165
5166
5167
5168
5169
5169
5170
5171
5172
5173
5174
5175
5176
5177
5178
5179
5179
5180
5181
5182
5183
5184
5185
5186
5187
5188
5189
5189
5189
5190
5191
5192
5193
5194
5195
5196
5197
5198
5199
5199
5199
5200
5201
5202
5203
5204
5205
5206
5207
5208
5209
5209
5210
5211
5212
5213
5214
5215
5216
5217
5218
5219
5219
5220
5221
5222
5223
5224
5225
5226
5227
5228
5229
5229
5230
5231
5232
5233
5234
5235
5236
5237
5238
5239
5239
5240
5241
5242
5243
5244
5245
5246
5247
5248
5249
5249
5250
5251
5252
5253
5254
5255
5256
5257
5258
5259
5259
5260
5261
5262
5263
5264
5265
5266
5267
5268
5269
5269
5270
5271
5272
5273
5274
5275
5276
5277
5278
5279
5279
5280
5281
5282
5283
5284
5285
5286
5287
5288
5289
5289
5289
5290
5291
5292
5293
5294
5295
5296
5297
5298
5299
5299
5299
5300
5301
5302
5303
5304
5305
5306
5307
5308
5309
5309
5310
5311
5312
5313
5314
5315
5316
5317
5318
5319
5319
5320
5321
5322
5323
5324
5325
5326
5327
5328
5329
5329
5330
5331
5332
5333
5334
5335
5336
5337
5338
5339
5339
5340
5341
5342
5343
5344
5345
5346
5347
5348
5349
5349
5350
5351
5352
5353
5354
5355
5356
5357
5358
5359
5359
5360
5361
5362
5363
5364
5365
5366
5367
5368
5369
5369
5370
5371
5372
5373
5374
5375
5376
5377
5378
5379
5379
5380
5381
5382
5383
5384
5385
5386
5387
5388
5389
5389
5389
5390
5391
5392
5393
5394
5395
5396
5397
5398
5399
5399
5399
5400
5401
5402
5403
5404
5405
5406
5407
5408
5409
5409
5410
5411
5412
5413
5414
5415
5415
5416
5417
5418
5419
5419
5420
5421
5422
5423
5424
5425
5426
5427
5428
5429
5429
5430
5431
5432
5433
5434
5435
5436
5437
5438
5439
5439
5440
5441
5442
5443
5444
5445
5446
5447
5448
5449
5449
5450
5451
5452
5453
5454
5455
5456
5457
5458
5459
5459
5460
5461
5462
5463
5464
5465
5466
5467
5468
5469
5469
5470
5471
5472
5473
5474
5475
5476
5477
5478
5479
5479
5480
5481
5482
5483
5484
5485
5486
5487
5488
5489
5489
5489
5490
5491
5492
5493
5494
5495
5496
5497
5498
5499
5499
5499
5500
5501
5502
5503
5504
5505
5506
5507
5508
5509
5509
5510
5511
5512
5513
5514
5515
5516
5517
5518
5519
5519
5520
5521
5522
5523
5524
5525
5526
5527
5528
5529
5529
5530
5531
5532
5533
5534
5535
5536
5537
5538
5539
5539
5540
5541
5542
5543
5544
5545
5546
5547
5548
5549
5549
5550
5551
5552
5553
5554
5555
5556
5557
5558
5559
5559
5560
5561
5562
5563
5564
5565
5566
5567
5568
5569
5569
5570
5571
5572
5573
5574
5575
5576
5577
5578
5579
5579
5580
5581
5582
5583
5584
5585
5586
5587
5588
5589
5589
5589
5590
5591
5592
5593
5594
5595
5596
5597
5598
5599
5599
5599
5600
5601
5602
5603
5604
5605
5606
5607
5608
5609
5609
5610
5611
5612
5613
5614
5615
5616
5617
5618
5619
5619
5620
5621
5622
5623
5624
5625
5626
5627
5628
5629
5629
5630
5631
5632
5633
5634
5635
5636
5637
5638
5639
5639
5640
5641
5642
5643
5644
5645
5646
5647
5648
5649
5649
5650
5651
5652
5653
5654
5655
5656
5657
5658
5659
5659
5660
5661
5662
5663
5664
5665
5666
5667
5668
5669
5669
5670
5671
5672
5673
5674
5675
5676
5677
5678
5679
5679
5680
5681
5682
5683
5684
5685
5686
5687
5688
5689
5689
5689
5690
5691
5692
5693
5694
5695
5696
5697
5698
5699
5699
5699
5700
5701
5702
5703
5704
5705
5706
5707
5708
5709
5709
5710
5711
5712
5713
5714
5715
5716
5717
5718
5719
5719
5720
5721
5722
5723
5724
5725
5726
5727
5728
5729
5729
5730
5731
5732
5733
5734
5735
5736
5737
5738
5739
5739
5740
5741
5742
5743
5744
5745
5746
5747
5748
5749
5749
5750
5751
5752
5753
5754
5755
5756
5757
5758
5759
5759
5760
5761
5762
5763
5764
5765
5766
5767
5768
5769
5769
5770
5771
5772
5773
5774
5775
5776
5777
5778
5779
5779
5780
5781
5782
5783
5784
5785
5786
5787
5788
5789
5789
5789
5790
5791
5792
5793
5794
5795
5796
5797
5798
5799
5799
5799
5800
5801
5802
5803
5804
5805
5806
5807
5808
5809
5809
5810
5811
5812
5813
5814
5815
5816
5817
5818
5819
5819
5820
5821
5822
5823
5824
5825
5826
5827
5828
5829
5829
5830
5831
5832
5833
5834
5835
5836
5837
5838
5839
5839
5840
5841
5842
5843
5844
5845
5846
5847
5848
5849
5849
5850
5851
5852
5853
5854
5855
5856
5857
5858
5859
5859
5860
5861
5862
5863
5864
5865
5866
5867
5868
5869
5869
5870
5871
5872
5873
5874
5875
5876
5877
5878
5879
5879
5880
5881
5882
5883
5884
5885
5886
5887
5888
5889
5889
5889
5890
5891
5892
5893
5894
5895
5896
5897
5898
5899
5899
5899
5900
5901
5902
5903
5904
5905
5906
5907
5908
5909
5909
5910
5911
5912
5913
5914
5915
5916
5917
5918
5919
5919
5920
5921
5922
5923
5924
5925
5926
5927
5928
5929
5929
5930
5931
5932
5933
5934
5935
5936
5937
5938
5939
5939
5940
5941
5942
5943
5944
5945
5946
5947
5948
5949
5949
5950
5951
5952
5953
5954
5955
5956
5957
5958
5959
5959
5960
5961
5962
5963
5964
5965
5966
5967
5968
5969
5969
5970
5971
5972
5973
5974
5975
5976
5977
5978
5979
5979
5980
5981
5982
5983
5984
5985
5986
5987
5988
5989
5989
5989
5990
5991
5992
5993
5994
5995
5996
5997
5998
5999
5999
5999
6000
6001
6002
6003
6004
6005
6006
6007
6008
6009
6009
6010
6011
6012
6013
6014
6015
6016
6017
6018
6019
6019
6020
6021
6022
6023
6024
6025
6026
6027
6028
6029
6029
6030
6031
6032
6033
6034
6035
6036
6037
6038
6039
6039
6040
6041
6042
6043
6044
6045
6046
6047
6048
6049
6049
6050
6051
6052
6053
6054
6055
6056
6057
6058
6059
6059
6060
6061
6062
6063
6064
6065
6066
6067
6068
6069
6069
6070
6071
6072
6073
6074
6075
6076
6077
6078
6079
6079
6080
6081
6082
6083
6084
6085
6086
6087
6088
6089
6089
6089
6090
6091
6092
6093
6094
6095
6096
6097
6098
6099
6099
6099
6100
6101
6102
6103
6104
6105
6106
6107
6108
6109
6109
6110
6111
6112
6113
6114
6115
6116
6117
6118
6119
6119
6120
6121
6122
6123
6124

```

```

4272     dladm_status_t status;
4273     dld_ioc_macprop_t *dip;
4274     uint16_t pvid;
4275
4276     if (flags != 0)
4277         return (DLADM_STATUS_NOTSUP);
4278     *perm_flags = MAC_PROP_PERM_RW;
4279     *val_cnt = 1;
4280     dip = i_dladm_buf_alloc_by_id(sizeof (uint16_t), linkid, MAC_PROP_PVID,
4281         0, &status);
4282     if (dip == NULL)
4283         return (status);
4284     status = i_dladm_macprop(handle, dip, B_FALSE);
4285     if (status == DLADM_STATUS_OK) {
4286         (void) memcpy(&pvid, dip->pr_val, sizeof (pvid));
4287         (void) snprintf(*prop_val, DLADM_PROP_VAL_MAX, "%u", pvid);
4288     } else {
4289         (void) strlcpy(*prop_val, "?", DLADM_PROP_VAL_MAX);
4290     }
4291     free(dip);
4292     return (status);
4293 }
4294 /* ARGSUSED */
4295 static dladm_status_t
4296 set_bridge_pvid(dladm_handle_t handle, prop_desc_t *pd, datalink_id_t linkid,
4297     val_desc_t *vdp, uint_t val_cnt, uint_t flags, datalink_media_t media)
4298 {
4299     dladm_status_t status;
4300     dld_ioc_macprop_t *dip;
4301     uint16_t pvid;
4302
4303     dip = i_dladm_buf_alloc_by_id(sizeof (uint16_t), linkid, MAC_PROP_PVID,
4304         0, &status);
4305     if (dip == NULL)
4306         return (status);
4307     pvid = vdp->vd_val;
4308     (void) memcpy(dip->pr_val, &pvid, sizeof (pvid));
4309     status = i_dladm_macprop(handle, dip, B_TRUE);
4310     free(dip);
4311     if (status != DLADM_STATUS_OK)
4312         return (status);
4313
4314     /* Tell the running daemon, if any */
4315     return (dladm_bridge_refresh(handle, linkid));
4316 }
4317 }
4318 /* ARGSUSED */
4319 static dladm_status_t
4320 check_bridge_pvid(dladm_handle_t handle, struct prop_desc *pd,
4321     datalink_id_t linkid, char **prop_val, uint_t *val_cntp, uint_t flags,
4322     val_desc_t **vdpp, datalink_media_t media)
4323 {
4324     char          *cp;
4325     uint_t        val_cnt = *val_cntp;
4326     val_desc_t   *vdp = *vdpp;
4327
4328     if (val_cnt != 1)
4329         return (DLADM_STATUS_BADVALCNT);
4330
4331     if (prop_val == NULL) {
4332         vdp->vd_val = 1;
4333     } else {
4334         errno = 0;
4335         vdp->vd_val = strtoul(prop_val[0], &cp, 0);
4336         if (errno != 0 || *cp != '\0')
4337

```

```

4338                                         return (DLADM_STATUS_BADVAL);
4339 }
4340
4341     return (vdp->vd_val > VLAN_ID_MAX ? DLADM_STATUS_BADVAL :
4342             DLADM_STATUS_OK);
4343 }
4344
4345 dladm_status_t
4346 i_dladm_wlan_param(dladm_handle_t handle, datalink_id_t linkid, void *buf,
4347     mac_prop_id_t cmd, size_t len, boolean_t set)
4348 {
4349     uint32_t      flags;
4350     dladm_status_t status;
4351     uint32_t      media;
4352     dld_ioc_macprop_t *dip;
4353     void          *dp;
4354
4355     if ((status = dladm_datalink_id2info(handle, linkid, &flags, NULL,
4356         &media, NULL, 0)) != DLADM_STATUS_OK) {
4357         return (status);
4358     }
4359
4360     if (media != DL_WIFI)
4361         return (DLADM_STATUS_BADARG);
4362
4363     if (!(flags & DLADM_OPT_ACTIVE))
4364         return (DLADM_STATUS_TEMPOONLY);
4365
4366     if (len == (MAX_BUF_LEN - WIFI_BUF_OFFSET))
4367         len = MAX_BUF_LEN - sizeof (dld_ioc_macprop_t) - 1;
4368
4369     dip = i_dladm_buf_alloc_by_id(len, linkid, cmd, 0, &status);
4370     if (dip == NULL)
4371         return (DLADM_STATUS_NOMEM);
4372
4373     dp = (uchar_t *)dip->pr_val;
4374     if (set)
4375         (void) memcpy(dp, buf, len);
4376
4377     status = i_dladm_macprop(handle, dip, set);
4378     if (status == DLADM_STATUS_OK) {
4379         if (!set)
4380             (void) memcpy(buf, dp, len);
4381     }
4382
4383     free(dip);
4384     return (status);
4385 }
4386
4387 dladm_status_t
4388 dladm_parse_link_props(char *str, dladm_arg_list_t **listp, boolean_t novalues)
4389 {
4390     return (dladm_parse_args(str, listp, novalues));
4391 }
4392
4393 /*
4394  * Retrieve the one link property from the database
4395  */
4396 /*ARGSUSED*/
4397 static int
4398 i_dladm_get_one_prop(dladm_handle_t handle, datalink_id_t linkid,
4399     const char *prop_name, void *arg)
4400 {
4401     dladm_arg_list_t      *proplist = arg;
4402     dladm_arg_info_t      *aip = NULL;

```

```

4404     aip = &proplist->al_info[proplist->al_count];
4405     /*
4406      * it is fine to point to prop_name since prop_name points to the
4407      * prop_table[n].pd_name.
4408     */
4409     aip->ai_name = prop_name;
4410
4411     (void) dladm_get_linkprop(handle, linkid, DLADM_PROP_VAL_PERSISTENT,
4412      prop_name, aip->ai_val, &aip->ai_count);
4413
4414     if (aip->ai_count != 0)
4415       proplist->al_count++;
4416
4417     return (DLADM_WALK_CONTINUE);
4418 }
4419
4420 /*
4421  * Retrieve all link properties for a link from the database and
4422  * return a property list.
4423  */
4424
4425 dladm_status_t
4426 dladm_link_get_proplist(dladm_handle_t handle, datalink_id_t linkid,
4427   dladm_arg_list_t **listp)
4428 {
4429   dladm_arg_list_t      *list;
4430   dladm_status_t         status = DLADM_STATUS_OK;
4431
4432   list = calloc(1, sizeof (dladm_arg_list_t));
4433   if (list == NULL)
4434     return (dladm_errno2status(errno));
4435
4436   status = dladm_walk_linkprop(handle, linkid, list,
4437     i_dladm_get_one_prop);
4438
4439   *listp = list;
4440   return (status);
4441 }
4442
4443 /*
4444  * Retrieve the named property from a proplist, check the value and
4445  * convert to a kernel structure.
4446  */
4447 static dladm_status_t
4448 i_dladm_link_proplist_extract_one(dladm_handle_t handle,
4449   dladm_arg_list_t *proplist, const char *name, uint_t flags, void *arg)
4450 {
4451   dladm_status_t          status;
4452   dladm_arg_info_t        *aip = NULL;
4453   int                     i, j;
4454
4455   /* Find named property in proplist */
4456   for (i = 0; i < proplist->al_count; i++) {
4457     aip = &proplist->al_info[i];
4458     if (strcasecmp(aip->ai_name, name) == 0)
4459       break;
4460   }
4461
4462   /* Property not in list */
4463   if (i == proplist->al_count)
4464     return (DLADM_STATUS_OK);
4465
4466   for (i = 0; i < DLADM_MAX_PROPS; i++) {
4467     prop_desc_t            *pdp = &prop_table[i];
4468     val_desc_t              *vdp;

```

```

4469     vdp = malloc(sizeof (val_desc_t) * aip->ai_count);
4470     if (vdp == NULL)
4471       return (DLADM_STATUS_NOMEM);
4472
4473     if (strcasecmp(aip->ai_name, pdp->pd_name) != 0)
4474       continue;
4475
4476     if (aip->ai_val == NULL)
4477       return (DLADM_STATUS_BADARG);
4478
4479   /* Check property value */
4480   if (pdp->pd_check != NULL) {
4481     status = pdp->pd_check(handle, pdp, 0, aip->ai_val,
4482      &(aip->ai_count), flags, &vdp, 0);
4483   } else {
4484     status = DLADM_STATUS_BADARG;
4485   }
4486
4487   if (status != DLADM_STATUS_OK)
4488     return (status);
4489
4490   for (j = 0; j < DLADM_MAX_RSRC_PROP; j++) {
4491     resource_prop_t *rpp = &rsrc_prop_table[j];
4492
4493     if (strcasecmp(aip->ai_name, rpp->rp_name) != 0)
4494       continue;
4495
4496   /* Extract kernel structure */
4497   if (rpp->rp_extract != NULL) {
4498     status = rpp->rp_extract(vdp,
4499      aip->ai_count, arg);
4500   } else {
4501     status = DLADM_STATUS_BADARG;
4502   }
4503   break;
4504
4505   if (status != DLADM_STATUS_OK)
4506     return (status);
4507
4508   break;
4509 }
4510
4511   return (status);
4512 }
4513 }
4514
4515 /*
4516  * Extract properties from a proplist and convert to mac_resource_props_t.
4517  */
4518 dladm_status_t
4519 dladm_link_proplist_extract(dladm_handle_t handle, dladm_arg_list_t *proplist,
4520   mac_resource_props_t *mrp, uint_t flags)
4521 {
4522   dladm_status_t          status;
4523   int                     i;
4524
4525   for (i = 0; i < DLADM_MAX_RSRC_PROP; i++) {
4526     status = i_dladm_link_proplist_extract_one(handle,
4527       proplist, rsrc_prop_table[i].rp_name, flags, mrp);
4528     if (status != DLADM_STATUS_OK)
4529       return (status);
4530   }
4531
4532 }
4533
4534 static const char *
4535 dladm_perm2str(uint_t perm, char *buf)

```

```

4536 {
4537     (void) snprintf(buf, DLADM_STRSIZE, "%c%c",
4538         ((perm & MAC_PROP_PERM_READ) != 0) ? 'r' : '-',
4539         ((perm & MAC_PROP_PERM_WRITE) != 0) ? 'w' : '-');
4540     return (buf);
4541 }
4543 dladm_status_t
4544 dladm_get_state(dladm_handle_t handle, datalink_id_t linkid,
4545     link_state_t *state)
4546 {
4547     uint_t                perms;
4548     return (i_dladm_get_public_prop(handle, linkid, "state", 0,
4549             &perms, state, sizeof (*state)));
4550 }
4553 boolean_t
4554 dladm_attr_is_linkprop(const char *name)
4555 {
4556     /* non-property attribute names */
4557     const char *nonprop[] = {
4558         /* dlmgmtd core attributes */
4559         "name",
4560         "class",
4561         "media",
4562         FPHYMAJ,
4563         FPHYINST,
4564         FDEVNAME,
4565
4566         /* other attributes for vlan, aggr, etc */
4567         DLADM_ATTR_NAMES
4568     };
4569     boolean_t    is_nonprop = B_FALSE;
4570     int          i;
4571
4572     for (i = 0; i < sizeof (nonprop) / sizeof (nonprop[0]); i++) {
4573         if (strcmp(name, nonprop[i]) == 0) {
4574             is_nonprop = B_TRUE;
4575             break;
4576         }
4577     }
4578     return (!is_nonprop);
4579 }
4582 dladm_status_t
4583 dladm_linkprop_is_set(dladm_handle_t handle, datalink_id_t linkid,
4584     dladm_prop_type_t type, const char *prop_name, boolean_t *is_set)
4585 {
4586     char            *buf, **propvals;
4587     uint_t          valcnt = DLADM_MAX_PROP_VALCNT;
4588     int              i;
4589     dladm_status_t  status = DLADM_STATUS_OK;
4590     size_t           bufsize;
4591
4592     *is_set = B_FALSE;
4593
4594     bufsize = (sizeof (char *) + DLADM_PROP_VAL_MAX) *
4595             DLADM_MAX_PROP_VALCNT;
4596     if ((buf = malloc(1, bufsize)) == NULL)
4597         return (DLADM_STATUS_NOMEM);
4598
4599     propvals = (char **)(void *)buf;
4600     for (i = 0; i < valcnt; i++)
4601         propvals[i] = buf +

```

```

4602                                         sizeof (char *) * DLADM_MAX_PROP_VALCNT +
4603                                         i * DLADM_PROP_VAL_MAX;
4604 }
4606     if (dladm_get_linkprop(handle, linkid, type, prop_name, propvals,
4607         &valcnt) != DLADM_STATUS_OK) {
4608         goto done;
4609     }
4611     /*
4612      * valcnt is always set to 1 by get_pool(), hence we need to check
4613      * for a non-null string to see if it is set. For protection and
4614      * allowed-ips, we can check either the *propval or the valcnt.
4615      */
4616     if ((strcmp(prop_name, "pool") == 0 ||
4617         strcmp(prop_name, "protection") == 0 ||
4618         strcmp(prop_name, "allowed-ips") == 0) &&
4619         (strlen(*propvals) != 0)) {
4620         *is_set = B_TRUE;
4621     } else if ((strcmp(prop_name, "cpus") == 0) && (valcnt != 0)) {
4622         *is_set = B_TRUE;
4623     } else if ((strcmp(prop_name, "softmac") == 0) && (valcnt != 0) &&
4624         (strcmp(propvals[0], "true") == 0)) {
4625         *is_set = B_TRUE;
4626     }
4628 done:
4629     if (buf != NULL)
4630         free(buf);
4631     return (status);
4632 }
4634 /* ARGSUSED */
4635 static dladm_status_t
4636 get_linkmode_prop(dladm_handle_t handle, prop_desc_t *pdp,
4637     datalink_id_t linkid, char **prop_val, uint_t *val_cnt,
4638     datalink_media_t media, uint_t flags, uint_t *perm_flags)
4639 {
4640     char            *s;
4641     uint32_t        v;
4642     dladm_status_t  status;
4643
4644     status = i_dladm_get_public_prop(handle, linkid, pdp->pd_name, flags,
4645             perm_flags, &v, sizeof (v));
4646     if (status != DLADM_STATUS_OK)
4647         return (status);
4648
4649     switch (v) {
4650     case DLADM_PART_CM_MODE:
4651         s = "cm";
4652         break;
4653     case DLADM_PART_UD_MODE:
4654         s = "ud";
4655         break;
4656     default:
4657         s = "";
4658         break;
4659     }
4660     (void) sprintf(prop_val[0], DLADM_STRSIZE, "%s", s);
4661     *val_cnt = 1;
4662     return (DLADM_STATUS_OK);
4663 }
4664 }
```

```
*****
110682 Thu Feb 20 18:59:05 2014
new/usr/src/man/man1m/dladm.1m
2553 mac address should be a dladm link property
*****
1 .\" te
2 .\" Copyright (c) 2008, Sun Microsystems, Inc. All Rights Reserved
3 .\" Sun Microsystems, Inc. gratefully acknowledges The Open Group for permission
4 .\" The Institute of Electrical and Electronics Engineers and The Open Group, ha
5 .\" are reprinted and reproduced in electronic form in the Sun OS Reference Manu
6 .\" and Electronics Engineers, Inc and The Open Group. In the event of any discr
7 .\" This notice shall appear on any product containing this material.
8 .\" The contents of this file are subject to the terms of the Common Development
9 .\" See the License for the specific language governing permissions and limitati
10 .\" fields enclosed by brackets "[]" replaced with your own identifying informat
11 .TH DLDADM 1M "Feb 22, 2014"
11 .TH DLDADM 1M "Sep 23, 2009"
12 .SH NAME
13 dladm \- administer data links
14 .SH SYNOPSIS
15 .LP
16 .nf
17 \fBdladm show-link\fR [\fB-P\fR] [\fB-s\fR [\fB-i\fR \fIinterval\fR]] [\fB-p\fR
18 \fBdladm rename-link\fR [\fB-R\fR \fIroot-dir\fR] \fIlink\fR \fInew-link\fR
19 .fi
21 .LP
22 .nf
23 \fBdladm delete-phys\fR \fIphys-link\fR
24 \fBdladm show-phys\fR [\fB-P\fR] [\fB-m\fR] [\fB-p\fR] \fB-o\fR \fIfield\fR[...
25 .fi
27 .LP
28 .nf
29 \fBdladm create-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-P\fR \fIpolicy
30 [\fB-T\fR \fItime\fR] [\fB-u\fR \fIaddress\fR] \fB-l\fR \fIether-link\fR [
31 \fBdladm modify-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-P\fR \fIpolicy
32 [\fB-T\fR \fItime\fR] [\fB-u\fR \fIaddress\fR] \fIaggr-link\fR
33 \fBdladm delete-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fIaggr-link\fR
34 \fBdladm add-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-l\fR \fIether-link
35 \fIaggr-link\fR
36 \fBdladm remove-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-l\fR \fIether-l
37 \fIaggr-link\fR
38 \fBdladm show-aggr\fR [\fB-PLX\fR] [\fB-s\fR [\fB-i\fR \fIinterval\fR]] [\fB-p
39 [\fIaggr-link\fR]
40 .fi
42 .LP
43 .nf
44 \fBdladm create-bridge\fR [\fB-P\fR \fIprotect\fR] [\fB-R\fR \fIroot-dir\fR] [\f
45 [\fB-m\fR \fImax-age\fR] [\fB-h\fR \fIhello-time\fR] [\fB-d\fR \fIforward-d
46 [\fB-l\fR \fIlink\fR...] \fIbridge-name\fR
47 .fi
49 .LP
50 .nf
51 \fBdladm modify-bridge\fR [\fB-P\fR \fIprotect\fR] [\fB-R\fR \fIroot-dir\fR] [\f
52 [\fB-m\fR \fImax-age\fR] [\fB-h\fR \fIhello-time\fR] [\fB-d\fR \fIforward-d
53 \fIbridge-name\fR
54 .fi
56 .LP
57 .nf
58 \fBdladm delete-bridge\fR [\fB-R\fR \fIroot-dir\fR] \fIbridge-name\fR
59 .fi
```

```
61 .LP
62 .nf
63 \fBdladm add-bridge\fR [\fB-R\fR \fIroot-dir\fR] \fB-l\fR \fIlink\fR [\fB-l\fR \
64 .fi
66 .LP
67 .nf
68 \fBdladm remove-bridge\fR [\fB-R\fR \fIroot-dir\fR] \fB-l\fR \fIlink\fR [\fB-l\f
69 .fi
71 .LP
72 .nf
73 \fBdladm show-bridge\fR [\fB-flt\fR] [\fB-s\fR [\fB-i\fR \fIinterval\fR]] [[\fB-
74 [\fIbridge-name\fR]
75 .fi
77 .LP
78 .nf
79 \fBdladm create-vlan\fR [\fB-ft\fR] [\fB-R\fR \fIroot-dir\fR] \fB-l\fR \fIether-
80 \fBdladm delete-vlan\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fIvlan-link\fR
81 \fBdladm show-vlan\fR [\fB-P\fR] [[\fB-p\fR] \fB-o\fR \fIfield\fR[...]] [\fIvla
82 .fi
84 .LP
85 .nf
86 \fBdladm scan-wifi\fR [[\fB-p\fR] \fB-o\fR \fIfield\fR[...]] [\fIwifi-link\fR]
87 \fBdladm connect-wifi\fR [\fB-e\fR \fIessid\fR] [\fB-i\fR \fIbssid\fR] [\fB-k\fR
88 [\fB-s\fR none | wep | wpa ] [\fB-a\fR open | shared] [\fB-b\fR bss | ibss]
89 [\fB-m\fR a | b | g] [\fB-T\fR \fItime\fR] [\fIwifi-link\fR]
90 \fBdladm disconnect-wifi\fR [\fB-a\fR] [\fIwifi-link\fR]
91 \fBdladm show-wifi\fR [[\fB-p\fR] \fB-o\fR \fIfield\fR[...]] [\fIwifi-link\fR]
92 .fi
94 .LP
95 .nf
96 \fBdladm show-ether\fR [\fB-x\fR] [[\fB-p\fR] \fB-o\fR \fIfield\fR[...]] [\fIet
97 .fi
99 .LP
100 .nf
101 \fBdladm set-linkprop\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-p\fR \fIprop\f
102 \fBdladm reset-linkprop\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-p\fR \fIpro
103 \fBdladm show-linkprop\fR [\fB-P\fR] [[\fB-c\fR] \fB-o\fR \fIfield\fR[...]] [\f
104 .fi
106 .LP
107 .nf
108 \fBdladm create-secobj\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-f\fR \fIfi
109 \fBdladm delete-secobj\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fIsecobj\fR[...
110 \fBdladm show-secobj\fR [\fB-P\fR] [[\fB-p\fR] \fB-o\fR \fIfield\fR[...]] [\fI
111 .fi
113 .LP
114 .nf
115 \fBdladm create-vnic\fR [\fB-t\fR] \fB-l\fR \fIlink\fR [\fB-R\fR \fIroot-dir\fR]
116 {factory \fB-n\fR \fIslot-identifier\fR} | {random [\fB-x\fR \fIprefix\fR]
117 [\fB-v\fR \fIlan-id\fR] [\fB-p\fR \fIprop\fR=\fIvalue\fR[...]] \fIvnic-li
118 \fBdladm delete-vnic\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fIvnic-link\fR
119 \fBdladm show-vnic\fR [\fB-pp\fR] [\fB-s\fR [\fB-i\fR \fIinterval\fR]] [\fB-o\fR
120 [\fB-l\fR \fIlink\fR] \fIvnic-link\fR]
121 .fi
123 .LP
124 .nf
125 \fBdladm create-etherstub\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fIetherstub\f
126 \fBdladm delete-etherstub\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fIetherstub\f
```

```

127 \fBdadm show-etherstub\fR [\fIetherstub\fR]
128 .fi
130 .LP
131 .nf
132 \fBdadm create-iptun\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-T\fR \fItypes\f
133 \fIiptun-link\fR
134 \fBdadm modify-iptun\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-s\fR \fItsrc\f
135 \fBdadm delete-iptun\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fIiptun-link\fR
136 \fBdadm show-iptun\fR [\fB-p\fR] [\fB-p\fR] \fB-o\fR \fIfield\fR[,...]] [\fIip
137 .fi

139 .LP
140 .nf
141 \fBdadm show-usage\fR [\fB-a\fR] \fB-f\fR \fIfilename\fR [\fB-p\fR \fIplotfile\f
142 [\fB-e\fR \fItime\fR] [\fIlink\fR]
143 .fi

145 .SH DESCRIPTION
146 .sp
147 .LP
148 The \fBdadm\fR command is used to administer data-links. A data-link is
149 represented in the system as a \fBSTREAMS DLPI\fR (v2) interface which can be
150 plumbed under protocol stacks such as \fBTCP/IP\fR. Each data-link relies on
151 either a single network device or an aggregation of devices to send packets to
152 or receive packets from a network.
153 .sp
154 .LP
155 Each \fBdadm\fR subcommand operates on one of the following objects:
156 .sp
157 .ne 2
158 .na
159 \fB\fBlink\fR\fR
160 .ad
161 .sp .6
162 .RS 4n
163 A datalink, identified by a name. In general, the name can use any alphanumeric
164 characters (or the underscore, \fB_\fR), but must start with an alphabetic
165 character and end with a number. A datalink name can be at most 31 characters,
166 and the ending number must be between 0 and 4294967294 (inclusive). The ending
167 number must not begin with a zero. Datalink names between 3 and 8 characters
168 are recommended.
169 .sp
170 Some subcommands operate only on certain types or classes of datalinks. For
171 those cases, the following object names are used:
172 .sp
173 .ne 2
174 .na
175 \fB\fBphys-link\fR\fR
176 .ad
177 .sp .6
178 .RS 4n
179 A physical datalink.
180 .RE

182 .sp
183 .ne 2
184 .na
185 \fB\fBvlan-link\fR\fR
186 .ad
187 .sp .6
188 .RS 4n
189 A VLAN datalink.
190 .RE

192 .sp

```

```

193 .ne 2
194 .na
195 \fB\fBaggr-link\fR\fR
196 .ad
197 .sp .6
198 .RS 4n
199 An aggregation datalink (or a key; see NOTES).
200 .RE

202 .sp
203 .ne 2
204 .na
205 \fB\fBether-link\fR\fR
206 .ad
207 .sp .6
208 .RS 4n
209 A physical Ethernet datalink.
210 .RE

212 .sp
213 .ne 2
214 .na
215 \fB\fBwifi-link\fR\fR
216 .ad
217 .sp .6
218 .RS 4n
219 A WiFi datalink.
220 .RE

222 .sp
223 .ne 2
224 .na
225 \fB\fBvnic-link\fR\fR
226 .ad
227 .sp .6
228 .RS 4n
229 A virtual network interface created on a link or an \fBetherstub\fR. It is a
230 pseudo device that can be treated as if it were an network interface card on a
231 machine.
232 .RE

234 .sp
235 .ne 2
236 .na
237 \fB\fBiptun-link\fR\fR
238 .ad
239 .sp .6
240 .RS 4n
241 An IP tunnel link.
242 .RE

244 .RE

246 .sp
247 .ne 2
248 .na
249 \fB\fBdev\fR\fR
250 .ad
251 .sp .6
252 .RS 4n
253 A network device, identified by concatenation of a driver name and an instance
254 number.
255 .RE

257 .sp
258 .ne 2

```

```

259 .na
260 \fB\fBetherstub\fR\fR
261 .ad
262 .sp .6
263 .RS 4n
264 An Ethernet stub can be used instead of a physical NIC to create VNICs. VNICs
265 created on an \fBetherstub\fR will appear to be connected through a virtual
266 switch, allowing complete virtual networks to be built without physical
267 hardware.
268 .RE

270 .sp
271 .ne 2
272 .na
273 \fB\fBbridge\fR\fR
274 .ad
275 .sp .6
276 .RS 4n
277 A bridge instance, identified by an administratively-chosen name. The name may
278 use any alphanumeric characters or the underscore, \fB_\fR, but must start and
279 end with an alphabetic character. A bridge name can be at most 31 characters.
280 The name \fBdefault\fR is reserved, as are all names starting with \fBSUNW\fR.
281 .sp
282 Note that appending a zero (\fB0\fR) to a bridge name produces a valid link
283 name, used for observability.
284 .RE

286 .sp
287 .ne 2
288 .na
289 \fB\fBsecobj\fR\fR
290 .ad
291 .sp .6
292 .RS 4n
293 A secure object, identified by an administratively-chosen name. The name can
294 use any alphanumeric characters, as well as underscore (\fB_\fR), period
295 (\fB.\fR), and hyphen (\fB-\fR). A secure object name can be at most 32
296 characters.
297 .RE

299 .SS "Options"
300 .sp
301 .LP
302 Each \fBdladm\fR subcommand has its own set of options. However, many of the
303 subcommands have the following as a common option:
304 .sp
305 .ne 2
306 .na
307 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
308 .ad
309 .sp .6
310 .RS 4n
311 Specifies an alternate root directory where the operation-such as creation,
312 deletion, or renaming-should apply.
313 .RE

315 .SS "SUBCOMMANDS"
316 .sp
317 .LP
318 The following subcommands are supported:
319 .sp
320 .ne 2
321 .na
322 \fB\fBdladm show-link\fR [\fB-P\fR] [\fB-s\fR [\fB-i\fR \fIinterval\fR]]
323 [\fB-p\fR \fB-o\fR \fIfield\fR[\fR[\fR[...]\fR]\fR]\fR
324 .ad

```

```

325 .sp .6
326 .RS 4n
327 Show link configuration information (the default) or statistics, either for all
328 datalinks or for the specified link \fIlink\fR. By default, the system is
329 configured with one datalink for each known network device.
330 .sp
331 .ne 2
332 .na
333 \fB\fB-o\fR \fIfield\fR[\fR[...]\fR], \fB--output\fR=\fIfield\fR[\fR[...]\fR]\fR
334 .ad
335 .sp .6
336 .RS 4n
337 A case-insensitive, comma-separated list of output fields to display. When not
338 modified by the \fB-s\fR option (described below), the field name must be one
339 of the fields listed below, or the special value \fBall\fR to display all
340 fields. By default (without \fB-o\fR), \fBshow-link\fR displays all fields.
341 .sp
342 .ne 2
343 .na
344 \fB\fBLINK\fR\fR
345 .ad
346 .sp .6
347 .RS 4n
348 The name of the datalink.
349 .RE

351 .sp
352 .ne 2
353 .na
354 \fB\fBCCLASS\fR\fR
355 .ad
356 .sp .6
357 .RS 4n
358 The class of the datalink. \fBdladm\fR distinguishes between the following
359 classes:
360 .sp
361 .ne 2
362 .na
363 \fB\fBphys\fR\fR
364 .ad
365 .sp .6
366 .RS 4n
367 A physical datalink. The \fBshow-phys\fR subcommand displays more detail for
368 this class of datalink.
369 .RE

371 .sp
372 .ne 2
373 .na
374 \fB\fBaggr\fR\fR
375 .ad
376 .sp .6
377 .RS 4n
378 An IEEE 802.3ad link aggregation. The \fBshow-aggr\fR subcommand displays more
379 detail for this class of datalink.
380 .RE

382 .sp
383 .ne 2
384 .na
385 \fB\fBvlan\fR\fR
386 .ad
387 .sp .6
388 .RS 4n
389 A VLAN datalink. The \fBshow-vlan\fR subcommand displays more detail for this
390 class of datalink.

```

```

391 .RE
393 .sp
394 .ne 2
395 .na
396 \fB\fBvnic\fR\fR
397 .ad
398 .sp .6
399 .RS 4n
400 A virtual network interface. The \fBshow-vnic\fR subcommand displays more
401 detail for this class of datalink.
402 .RE

404 .RE

406 .sp
407 .ne 2
408 .na
409 \fB\fBMTU\fR\fR
410 .ad
411 .sp .6
412 .RS 4n
413 The maximum transmission unit size for the datalink being displayed.
414 .RE

416 .sp
417 .ne 2
418 .na
419 \fB\fBSTATE\fR\fR
420 .ad
421 .sp .6
422 .RS 4n
423 The link state of the datalink. The state can be \fBup\fR, \fBdown\fR, or
424 \fBunknown\fR.
425 .RE

427 .sp
428 .ne 2
429 .na
430 \fB\fBBRIDGE\fR\fR
431 .ad
432 .sp .6
433 .RS 4n
434 The name of the bridge to which this link is assigned, if any.
435 .RE

437 .sp
438 .ne 2
439 .na
440 \fB\fBOVER\fR\fR
441 .ad
442 .sp .6
443 .RS 4n
444 The physical datalink(s) over which the datalink is operating. This applies to
445 \fBaggr\fR, \fBbridge\fR, and \fBvlan\fR classes of datalinks. A VLAN is
446 created over a single physical datalink, a bridge has multiple attached links,
447 and an aggregation is comprised of one or more physical datalinks.
448 .RE

450 When the \fB-o\fR option is used in conjunction with the \fB-s\fR option, used
451 to display link statistics, the field name must be one of the fields listed
452 below, or the special value \fBall\fR to display all fields
453 .sp
454 .ne 2
455 .na
456 \fB\fBLINK\fR\fR

```

```

457 .ad
458 .sp .6
459 .RS 4n
460 The name of the datalink.
461 .RE

463 .sp
464 .ne 2
465 .na
466 \fB\fBIPACKETS\fR\fR
467 .ad
468 .sp .6
469 .RS 4n
470 Number of packets received on this link.
471 .RE

473 .sp
474 .ne 2
475 .na
476 \fB\fBRBYTES\fR\fR
477 .ad
478 .sp .6
479 .RS 4n
480 Number of bytes received on this link.
481 .RE

483 .sp
484 .ne 2
485 .na
486 \fB\fBIERRORS\fR\fR
487 .ad
488 .sp .6
489 .RS 4n
490 Number of input errors.
491 .RE

493 .sp
494 .ne 2
495 .na
496 \fB\fBOPACKETS\fR\fR
497 .ad
498 .sp .6
499 .RS 4n
500 Number of packets sent on this link.
501 .RE

503 .sp
504 .ne 2
505 .na
506 \fB\fBOBYTES\fR\fR
507 .ad
508 .sp .6
509 .RS 4n
510 Number of bytes received on this link.
511 .RE

513 .sp
514 .ne 2
515 .na
516 \fB\fBOERRORS\fR\fR
517 .ad
518 .sp .6
519 .RS 4n
520 Number of output errors.
521 .RE

```

```

523 .RE
525 .sp
526 .ne 2
527 .na
528 \fB\fB-p\fR, \fB--parseable\fR\fR
529 .ad
530 .sp .6
531 .RS 4n
532 Display using a stable machine-parseable format. The \fB-o\fR option is
533 required with \fB-p\fR. See "Parseable Output Format", below.
534 .RE

536 .sp
537 .ne 2
538 .na
539 \fB\fB-P\fR, \fB--persistent\fR\fR
540 .ad
541 .sp .6
542 .RS 4n
543 Display the persistent link configuration.
544 .RE

546 .sp
547 .ne 2
548 .na
549 \fB\fB-s\fR, \fB--statistics\fR\fR
550 .ad
551 .sp .6
552 .RS 4n
553 Display link statistics.
554 .RE

556 .sp
557 .ne 2
558 .na
559 \fB\fB-i\fR \fIinterval\fR, \fB--interval\fR=\fIinterval\fR\fR
560 .ad
561 .sp .6
562 .RS 4n
563 Used with the \fB-s\fR option to specify an interval, in seconds, at which
564 statistics should be displayed. If this option is not specified, statistics
565 will be displayed only once.
566 .RE

568 .RE

570 .sp
571 .ne 2
572 .na
573 \fB\fBdadm rename-link\fR [\fB-R\fR \fIroot-dir\fR] \fIlink\fR
574 \fInew-link\fR\fR
575 .ad
576 .sp .6
577 .RS 4n
578 Rename \fIlink\fR to \fInew-link\fR. This is used to give a link a meaningful
579 name, or to associate existing link configuration such as link properties of a
580 removed device with a new device. See the \fBEXAMPLES\fR section for specific
581 examples of how this subcommand is used.
582 .sp
583 .ne 2
584 .na
585 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
586 .ad
587 .sp .6
588 .RS 4n

```

```

589 See "Options," above.
590 .RE
592 .RE
594 .sp
595 .ne 2
596 .na
597 \fB\fBdadm delete-phys\fR \fIphys-link\fR\fR
598 .ad
599 .sp .6
600 .RS 4n
601 This command is used to delete the persistent configuration of a link
602 associated with physical hardware which has been removed from the system. See
603 the \fBEXAMPLES\fR section.
604 .RE

606 .sp
607 .ne 2
608 .na
609 \fB\fBdadm show-phys\fR [\fB-P\fR] [[\fB-p\fR] \fB-o\fR \fIfield\fR[,...]]
610 [\fB-H\fR] [\fIphys-link\fR]\fR
611 .ad
612 .sp .6
613 .RS 4n
614 Show the physical device and attributes of all physical links, or of the named
615 physical link. Without \fB-P\fR, only physical links that are available on the
616 running system are displayed.
617 .sp
618 .ne 2
619 .na
620 \fB\fB-H\fR\fR
621 .ad
622 .sp .6
623 .RS 4n
624 Show hardware resource usage, as returned by the NIC driver. Output from
625 \fB-H\fR displays the following elements:
626 .sp
627 .ne 2
628 .na
629 \fB\fBLINK\fR\fR
630 .ad
631 .sp .6
632 .RS 4n
633 A physical device corresponding to a NIC driver.
634 .RE

636 .sp
637 .ne 2
638 .na
639 \fB\fBGROUP\fR\fR
640 .ad
641 .sp .6
642 .RS 4n
643 A collection of rings.
644 .RE

646 .sp
647 .ne 2
648 .na
649 \fB\fBGROUPTYPE\fR\fR
650 .ad
651 .sp .6
652 .RS 4n
653 RX or TX. All rings in a group are of the same group type.
654 .RE

```

```

656 .sp
657 .ne 2
658 .na
659 \fB\fBRINGS\fR\fR
660 .ad
661 .sp .6
662 .RS 4n
663 A hardware resource used by a data link, subject to assignment by a driver to
664 different groups.
665 .RE

667 .sp
668 .ne 2
669 .na
670 \fB\fBCLIENTS\fR\fR
671 .ad
672 .sp .6
673 .RS 4n
674 MAC clients that are using the rings within a group.
675 .RE

677 .RE

679 .sp
680 .ne 2
681 .na
682 \fB\fB-o\fR \fIfield\fR, \fB--output\fR=\fIfield\fR\fR
683 .ad
684 .sp .6
685 .RS 4n
686 A case-insensitive, comma-separated list of output fields to display. The field
687 name must be one of the fields listed below, or the special value \fBall\fR, to
688 display all fields. For each link, the following fields can be displayed:
689 .sp
690 .ne 2
691 .na
692 \fB\fBLINK\fR\fR
693 .ad
694 .sp .6
695 .RS 4n
696 The name of the datalink.
697 .RE

699 .sp
700 .ne 2
701 .na
702 \fB\fBMEDIA\fR\fR
703 .ad
704 .sp .6
705 .RS 4n
706 The media type provided by the physical datalink.
707 .RE

709 .sp
710 .ne 2
711 .na
712 \fB\fBSTATE\fR\fR
713 .ad
714 .sp .6
715 .RS 4n
716 The state of the link. This can be \fBup\fR, \fBdown\fR, or \fBunknown\fR.
717 .RE

719 .sp
720 .ne 2

```

```

721 .na
722 \fB\fBSPEED\fR\fR
723 .ad
724 .sp .6
725 .RS 4n
726 The current speed of the link, in megabits per second.
727 .RE

729 .sp
730 .ne 2
731 .na
732 \fB\fBDUPLEX\fR\fR
733 .ad
734 .sp .6
735 .RS 4n
736 For Ethernet links, the full/half duplex status of the link is displayed if the
737 link state is \fBup\fR. The duplex is displayed as \fBUnknown\fR in all other
738 cases.
739 .RE

741 .sp
742 .ne 2
743 .na
744 \fB\fBDEVICE\fR\fR
745 .ad
746 .sp .6
747 .RS 4n
748 The name of the physical device under this link.
749 .RE

751 .RE

753 .sp
754 .ne 2
755 .na
756 \fB\fB-p\fR, \fB--parseable\fR\fR
757 .ad
758 .sp .6
759 .RS 4n
760 Display using a stable machine-parseable format. The \fB-o\fR option is
761 required with \fB-p\fR. See "Parseable Output Format", below.
762 .RE

764 .sp
765 .ne 2
766 .na
767 \fB\fB-P\fR, \fB--persistent\fR\fR
768 .ad
769 .sp .6
770 .RS 4n
771 This option displays persistent configuration for all links, including those
772 that have been removed from the system. The output provides a \fBFLAGS\fR
773 column in which the \fBr\fR flag indicates that the physical device associated
774 with a physical link has been removed. For such links, \fBdelete-phys\fR can be
775 used to purge the link's configuration from the system.
776 .RE

778 .RE

780 .sp
781 .ne 2
782 .na
783 \fB\fBdadm create-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-P\fR
784 \fIpolicy\fR] [\fB-L\fR \fImode\fR] [\fB-T\fR \fItime\fR] [\fB-u\fR
785 \fIaddress\fR] \fB-l\fR \fIether-link1\fR [\fB-l\fR \fIether-link2\fR...]
786 \fIaggr-link\fR\fR

```

```

787 .ad
788 .sp .6
789 .RS 4n
790 Combine a set of links into a single IEEE 802.3ad link aggregation named
791 \fIaggr-link\fR. The use of an integer \fIkey\fR to generate a link name for
792 the aggregation is also supported for backward compatibility. Many of the
793 \fB*\fR\fB-aggr\fR subcommands below also support the use of a \fIkey\fR to
794 refer to a given aggregation, but use of the aggregation link name is
795 preferred. See the \fBNOTES\fR section for more information on keys.
796 .sp
797 \fBdladm\fR supports a number of port selection policies for an aggregation of
798 ports. (See the description of the \fB-P\fR option, below.) If you do not
799 specify a policy, \fBcreate-aggr\fR uses the default, the L4 policy, described
800 under the \fB-P\fR option.
801 .sp
802 .ne 2
803 .na
804 \fB\fB-1\fR \fIether-link\fR, \fB--link\fR=\fIether-link\fR\fR
805 .ad
806 .sp .6
807 .RS 4n
808 Each Ethernet link (or port) in the aggregation is specified using an \fB-1\fR
809 option followed by the name of the link to be included in the aggregation.
810 Multiple links are included in the aggregation by specifying multiple \fB-1\fR
811 options. For backward compatibility with previous versions of Solaris, the
812 \fBdladm\fR command also supports the using the \fB-d\fR option (or
813 \fB--dev\fR) with a device name to specify links by their underlying device
814 name. The other \fB*\fR\fB-aggr\fR subcommands that take \fB-1\fRoptions also
815 accept \fB-d\fR.
816 .RE

818 .sp
819 .ne 2
820 .na
821 \fB\fB-t\fR, \fB--temporary\fR\fR
822 .ad
823 .sp .6
824 .RS 4n
825 Specifies that the aggregation is temporary. Temporary aggregations last until
826 the next reboot.
827 .RE

829 .sp
830 .ne 2
831 .na
832 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
833 .ad
834 .sp .6
835 .RS 4n
836 See "Options," above.
837 .RE

839 .sp
840 .ne 2
841 .na
842 \fB\fB-P\fR \fIpolicy\fR, \fB--policy\fR=\fIpolicy\fR\fR
843 .ad
844 .br
845 .na
846 \fB\fR
847 .ad
848 .sp .6
849 .RS 4n
850 Specifies the port selection policy to use for load spreading of outbound
851 traffic. The policy specifies which \fIdev\fR object is used to send packets. A
852 policy is a list of one or more layers specifiers separated by commas. A layer

```

```

853 specifier is one of the following:
854 .sp
855 .ne 2
856 .na
857 \fB\fBL2\fR\fR
858 .ad
859 .sp .6
860 .RS 4n
861 Select outbound device according to source and destination \fBMAC\fR addresses
862 of the packet.
863 .RE

865 .sp
866 .ne 2
867 .na
868 \fB\fBL3\fR\fR
869 .ad
870 .sp .6
871 .RS 4n
872 Select outbound device according to source and destination \fBIP\fR addresses
873 of the packet.
874 .RE

876 .sp
877 .ne 2
878 .na
879 \fB\fBL4\fR\fR
880 .ad
881 .sp .6
882 .RS 4n
883 Select outbound device according to the upper layer protocol information
884 contained in the packet. For \fBTCP\fR and \fBUDP\fR, this includes source and
885 destination ports. For IPsec, this includes the \fBSPI\fR (Security Parameters
886 Index).
887 .RE

889 For example, to use upper layer protocol information, the following policy can
890 be used:
891 .sp
892 .in +2
893 .nf
894 -P L4
895 .fi
896 .in -2
897 .sp

899 Note that policy L4 is the default.
900 .sp
901 To use the source and destination \fBMAC\fR addresses as well as the source and
902 destination \fBIP\fR addresses, the following policy can be used:
903 .sp
904 .in +2
905 .nf
906 -P L2,L3
907 .fi
908 .in -2
909 .sp

911 .RE

913 .sp
914 .ne 2
915 .na
916 \fB\fBL\fR \fImode\fR, \fB--lacp-mode\fR=\fImode\fR\fR
917 .ad
918 .sp .6

```

```

919 .RS 4n
920 Specifies whether \fBLACP\fR should be used and, if used, the mode in which it
921 should operate. Supported values are \fBoff\fR, \fBactive\fR or \fBpassive\fR.
922 .RE

924 .sp
925 .ne 2
926 .na
927 \fB\fB-T\fR \fItimer\fR, \fB--lacp-timer\fR=\fItimer\fR\fR\fR
928 .ad
929 .br
930 .na
931 \fB\fR
932 .ad
933 .sp .6
934 .RS 4n
935 Specifies the \fBLACP\fR timer value. The supported values are \fBshort\fR or
936 \fBlong\fRjjj.
937 .RE

939 .sp
940 .ne 2
941 .na
942 \fB\fB-u\fR \fIaddress\fR, \fB--unicast\fR=\fIaddress\fR\fR\fR
943 .ad
944 .sp .6
945 .RS 4n
946 Specifies a fixed unicast hardware address to be used for the aggregation. If
947 this option is not specified, then an address is automatically chosen from the
948 set of addresses of the component devices.
949 .RE

951 .RE

953 .sp
954 .ne 2
955 .na
956 \fB\fBdladm modify-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-P\fR
957 \fIpolicy\fR] [\fB-L\fR \fImode\fR] [\fB-T\fR \fItimer\fR] [\fB-u\fR
958 \fIaddress\fR] \fIaggr-link\fR\fR
959 .ad
960 .sp .6
961 .RS 4n
962 Modify the parameters of the specified aggregation.
963 .sp
964 .ne 2
965 .na
966 \fB\fB-t\fR, \fB--temporary\fR\fR
967 .ad
968 .sp .6
969 .RS 4n
970 Specifies that the modification is temporary. Temporary aggregations last until
971 the next reboot.
972 .RE

974 .sp
975 .ne 2
976 .na
977 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR\fR
978 .ad
979 .sp .6
980 .RS 4n
981 See "Options," above.
982 .RE

984 .sp

```

```

985 .ne 2
986 .na
987 \fB\fB-P\fR \fIpolicy\fR, \fB--policy\fR=\fIpolicy\fR\fR\fR
988 .ad
989 .sp .6
990 .RS 4n
991 Specifies the port selection policy to use for load spreading of outbound
992 traffic. See \fBdladm create-aggr\fR for a description of valid policy values.
993 .RE

995 .sp
996 .ne 2
997 .na
998 \fB\fB-L\fR \fImode\fR, \fB--lacp-mode\fR=\fImode\fR\fR\fR
999 .ad
1000 .sp .6
1001 .RS 4n
1002 Specifies whether \fBLACP\fR should be used and, if used, the mode in which it
1003 should operate. Supported values are \fBoff\fR, \fBactive\fR, or \fBpassive\fR.
1004 .RE

1006 .sp
1007 .ne 2
1008 .na
1009 \fB\fB-T\fR \fItimer\fR, \fB--lacp-timer\fR=\fItimer\fR\fR\fR
1010 .ad
1011 .br
1012 .na
1013 \fB\fR
1014 .ad
1015 .sp .6
1016 .RS 4n
1017 Specifies the \fBLACP\fR timer value. The supported values are \fBshort\fR or
1018 \fBlong\fR.
1019 .RE

1021 .sp
1022 .ne 2
1023 .na
1024 \fB\fB-u\fR \fIaddress\fR, \fB--unicast\fR=\fIaddress\fR\fR\fR
1025 .ad
1026 .sp .6
1027 .RS 4n
1028 Specifies a fixed unicast hardware address to be used for the aggregation. If
1029 this option is not specified, then an address is automatically chosen from the
1030 set of addresses of the component devices.
1031 .RE

1033 .RE

1035 .sp
1036 .ne 2
1037 .na
1038 \fB\fBdladm delete-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR]
1039 \fIaggr-link\fR\fR
1040 .ad
1041 .sp .6
1042 .RS 4n
1043 Deletes the specified aggregation.
1044 .sp
1045 .ne 2
1046 .na
1047 \fB\fB-t\fR, \fB--temporary\fR\fR
1048 .ad
1049 .sp .6
1050 .RS 4n

```

```

1051 Specifies that the deletion is temporary. Temporary deletions last until the
1052 next reboot.
1053 .RE

1055 .sp
1056 .ne 2
1057 .na
1058 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
1059 .ad
1060 .sp .6
1061 .RS 4n
1062 See "Options," above.
1063 .RE

1065 .RE

1067 .sp
1068 .ne 2
1069 .na
1070 \fB\fBdadm add-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-l\fR
1071 \fIether-link1\fR [\fB--link\fR=\fIether-link2\fR...] \fIaggr-link\fR\fR
1072 .ad
1073 .sp .6
1074 .RS 4n
1075 Adds links to the specified aggregation.
1076 .sp
1077 .ne 2
1078 .na
1079 \fB\fB-l\fR \fIether-link\fR, \fB--link\fR=\fIether-link\fR\fR
1080 .ad
1081 .sp .6
1082 .RS 4n
1083 Specifies an Ethernet link to add to the aggregation. Multiple links can be
1084 added by supplying multiple \fB-l\fR options.
1085 .RE

1087 .sp
1088 .ne 2
1089 .na
1090 \fB\fB-t\fR, \fB--temporary\fR\fR
1091 .ad
1092 .sp .6
1093 .RS 4n
1094 Specifies that the additions are temporary. Temporary additions last until the
1095 next reboot.
1096 .RE

1098 .sp
1099 .ne 2
1100 .na
1101 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
1102 .ad
1103 .sp .6
1104 .RS 4n
1105 See "Options," above.
1106 .RE

1108 .RE

1110 .sp
1111 .ne 2
1112 .na
1113 \fB\fBdadm remove-aggr\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-l\fR
1114 \fIether-link1\fR [\fB--link\fR=\fIether-link2\fR...] \fIaggr-link\fR\fR
1115 .ad
1116 .sp .6

```

```

1117 .RS 4n
1118 Removes links from the specified aggregation.
1119 .sp
1120 .ne 2
1121 .na
1122 \fB\fB-l\fR \fIether-link\fR, \fB--link\fR=\fIether-link\fR\fR
1123 .ad
1124 .sp .6
1125 .RS 4n
1126 Specifies an Ethernet link to remove from the aggregation. Multiple links can
1127 be added by supplying multiple \fB-l\fR options.
1128 .RE

1130 .sp
1131 .ne 2
1132 .na
1133 \fB\fB-t\fR, \fB--temporary\fR\fR
1134 .ad
1135 .sp .6
1136 .RS 4n
1137 Specifies that the removals are temporary. Temporary removal last until the
1138 next reboot.
1139 .RE

1141 .sp
1142 .ne 2
1143 .na
1144 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
1145 .ad
1146 .sp .6
1147 .RS 4n
1148 See "Options," above.
1149 .RE

1151 .RE

1153 .sp
1154 .ne 2
1155 .na
1156 \fB\fBdadm show-aggr\fR [\fB-PLx\fR] [\fB-s\fR [\fB-i\fR \fIinterval\fR]]
1157 [\fB-p\fR] \fB-o\fR \fIfield\fR[,...] [\fIaggr-link\fR]\fR
1158 .ad
1159 .sp .6
1160 .RS 4n
1161 Show aggregation configuration (the default), \fBLACP\fR information, or
1162 statistics, either for all aggregations or for the specified aggregation.
1163 .sp
1164 By default (with no options), the following fields can be displayed:
1165 .sp
1166 .ne 2
1167 .na
1168 \fB\fBLINK\fR\fR
1169 .ad
1170 .sp .6
1171 .RS 4n
1172 The name of the aggregation link.
1173 .RE

1175 .sp
1176 .ne 2
1177 .na
1178 \fB\fBPOLICY\fR\fR
1179 .ad
1180 .sp .6
1181 .RS 4n
1182 The LACP policy of the aggregation. See the \fBcreate-aggr\fR \fB-P\fR option

```

```

1183 for a description of the possible values.
1184 .RE

1186 .sp
1187 .ne 2
1188 .na
1189 \fB\fBADDRPOLICY\fR\fR
1190 .ad
1191 .sp .6
1192 .RS 4n
1193 Either \fBauto\fR, if the aggregation is configured to automatically configure
1194 its unicast MAC address (the default if the \fB-u\fR option was not used to
1195 create or modify the aggregation), or \fBfixed\fR, if \fB-u\fR was used to set
1196 a fixed MAC address.
1197 .RE

1199 .sp
1200 .ne 2
1201 .na
1202 \fB\fBLACPACTIVITY\fR\fR
1203 .ad
1204 .sp .6
1205 .RS 4n
1206 The LACP mode of the aggregation. Possible values are \fBoff\fR, \fBactive\fR,
1207 or \fBpassive\fR, as set by the \fB-l\fR option to \fBcreate-aggr\fR or
1208 \fBmodify-aggr\fR.
1209 .RE

1211 .sp
1212 .ne 2
1213 .na
1214 \fB\fBLACPTIMER\fR\fR
1215 .ad
1216 .sp .6
1217 .RS 4n
1218 The LACP timer value of the aggregation as set by the \fB-T\fR option of
1219 \fBcreate-aggr\fR or \fBmodify-aggr\fR.
1220 .RE

1222 .sp
1223 .ne 2
1224 .na
1225 \fB\fBFLAGS\fR\fR
1226 .ad
1227 .sp .6
1228 .RS 4n
1229 A set of state flags associated with the aggregation. The only possible flag is
1230 \fBF\fR, which is displayed if the administrator forced the creation the
1231 aggregation using the \fB-f\fR option to \fBcreate-aggr\fR. Other flags might
1232 be defined in the future.
1233 .RE

1235 The \fBshow-aggr\fR command accepts the following options:
1236 .sp
1237 .ne 2
1238 .na
1239 \fB\fBLACPL\fR, \fB--lacp\fR\fR
1240 .ad
1241 .sp .6
1242 .RS 4n
1243 Displays detailed \fBLACP\fR information for the aggregation link and each
1244 underlying port. Most of the state information displayed by this option is
1245 defined by IEEE 802.3. With this option, the following fields can be displayed:
1246 .sp
1247 .ne 2
1248 .na

```

```

1249 \fB\fBLINK\fR\fR
1250 .ad
1251 .sp .6
1252 .RS 4n
1253 The name of the aggregation link.
1254 .RE

1256 .sp
1257 .ne 2
1258 .na
1259 \fB\fBPORt\fR\fR
1260 .ad
1261 .sp .6
1262 .RS 4n
1263 The name of one of the underlying aggregation ports.
1264 .RE

1266 .sp
1267 .ne 2
1268 .na
1269 \fB\fBAGGREGATABLE\fR\fR
1270 .ad
1271 .sp .6
1272 .RS 4n
1273 Whether the port can be added to the aggregation.
1274 .RE

1276 .sp
1277 .ne 2
1278 .na
1279 \fB\fBSYNC\fR\fR
1280 .ad
1281 .sp .6
1282 .RS 4n
1283 If \fByes\fR, the system considers the port to be synchronized and part of the
1284 aggregation.
1285 .RE

1287 .sp
1288 .ne 2
1289 .na
1290 \fB\fBCOLL\fR\fR
1291 .ad
1292 .sp .6
1293 .RS 4n
1294 If \fByes\fR, collection of incoming frames is enabled on the associated port.
1295 .RE

1297 .sp
1298 .ne 2
1299 .na
1300 \fB\fBDIST\fR\fR
1301 .ad
1302 .sp .6
1303 .RS 4n
1304 If \fByes\fR, distribution of outgoing frames is enabled on the associated
1305 port.
1306 .RE

1308 .sp
1309 .ne 2
1310 .na
1311 \fB\fBDEFAULTED\fR\fR
1312 .ad
1313 .sp .6
1314 .RS 4n

```

1315 If \fByes\fR, the port is using defaulted partner information (that is, has not  
 1316 received LACP data from the LACP partner).  
 1317 .RE

1319 .sp  
 1320 .ne 2  
 1321 .na  
 1322 \fB\fBEXPIRED\fR\fR  
 1323 .ad  
 1324 .sp .6  
 1325 .RS 4n  
 1326 If \fByes\fR, the receive state of the port is in the \fBEXPIRED\fR state.  
 1327 .RE

1329 .RE

1331 .sp  
 1332 .ne 2  
 1333 .na  
 1334 \fB\fB-x\fR, \fB--extended\fR\fR  
 1335 .ad  
 1336 .sp .6  
 1337 .RS 4n  
 1338 Display additional aggregation information including detailed information on  
 1339 each underlying port. With \fB-x\fR, the following fields can be displayed:  
 1340 .sp  
 1341 .ne 2  
 1342 .na  
 1343 \fB\fBLINK\fR\fR  
 1344 .ad  
 1345 .sp .6  
 1346 .RS 4n  
 1347 The name of the aggregation link.  
 1348 .RE

1350 .sp  
 1351 .ne 2  
 1352 .na  
 1353 \fB\fBPORt\fR\fR  
 1354 .ad  
 1355 .sp .6  
 1356 .RS 4n  
 1357 The name of one of the underlying aggregation ports.  
 1358 .RE

1360 .sp  
 1361 .ne 2  
 1362 .na  
 1363 \fB\fBSPEED\fR\fR  
 1364 .ad  
 1365 .sp .6  
 1366 .RS 4n  
 1367 The speed of the link or port in megabits per second.  
 1368 .RE

1370 .sp  
 1371 .ne 2  
 1372 .na  
 1373 \fB\fBDUPLEX\fR\fR  
 1374 .ad  
 1375 .sp .6  
 1376 .RS 4n  
 1377 The full/half duplex status of the link or port is displayed if the link state  
 1378 is \fBup\fR. The duplex status is displayed as \fBunknown\fR in all other  
 1379 cases.  
 1380 .RE

1382 .sp  
 1383 .ne 2  
 1384 .na  
 1385 \fB\fBSTATE\fR\fR  
 1386 .ad  
 1387 .sp .6  
 1388 .RS 4n  
 1389 The link state. This can be \fBup\fR, \fBdown\fR, or \fBunknown\fR.  
 1390 .RE

1392 .sp  
 1393 .ne 2  
 1394 .na  
 1395 \fB\fBADDRESS\fR\fR  
 1396 .ad  
 1397 .sp .6  
 1398 .RS 4n  
 1399 The MAC address of the link or port.  
 1400 .RE

1402 .sp  
 1403 .ne 2  
 1404 .na  
 1405 \fB\fBPORtSTATE\fR\fR  
 1406 .ad  
 1407 .sp .6  
 1408 .RS 4n  
 1409 This indicates whether the individual aggregation port is in the \fBstandby\fR  
 1410 or \fBattached\fR state.  
 1411 .RE

1413 .RE

1415 .sp  
 1416 .ne 2  
 1417 .na  
 1418 \fB\fB-o\fR \fIfield\fR[...], \fB--output\fR=\fIfield\fR[...]\fR  
 1419 .ad  
 1420 .sp .6  
 1421 .RS 4n  
 1422 A case-insensitive, comma-separated list of output fields to display. The field  
 1423 name must be one of the fields listed above, or the special value \fBall\fR, to  
 1424 display all fields. The fields applicable to the \fB-o\fR option are limited to  
 1425 those listed under each output mode. For example, if using \fB-L\fR, only the  
 1426 fields listed under \fB-L\fR, above, can be used with \fB-o\fR.  
 1427 .RE

1429 .sp  
 1430 .ne 2  
 1431 .na  
 1432 \fB\fB-p\fR, \fB--parseable\fR\fR  
 1433 .ad  
 1434 .sp .6  
 1435 .RS 4n  
 1436 Display using a stable machine-parseable format. The \fB-o\fR option is  
 1437 required with \fB-p\fR. See "Parseable Output Format", below.  
 1438 .RE

1440 .sp  
 1441 .ne 2  
 1442 .na  
 1443 \fB\fB-P\fR, \fB--persistent\fR\fR  
 1444 .ad  
 1445 .sp .6  
 1446 .RS 4n

```

1447 Display the persistent aggregation configuration rather than the state of the
1448 running system.
1449 .RE

1451 .sp
1452 .ne 2
1453 .na
1454 \fB\fB-s\fR, \fB--statistics\fR\fR
1455 .ad
1456 .sp .6
1457 .RS 4n
1458 Displays aggregation statistics.
1459 .RE

1461 .sp
1462 .ne 2
1463 .na
1464 \fB\fB-i\fR \fIinterval\fR, \fB--interval\fR=\fIinterval\fR\fR
1465 .ad
1466 .sp .6
1467 .RS 4n
1468 Used with the \fB-s\fR option to specify an interval, in seconds, at which
1469 statistics should be displayed. If this option is not specified, statistics
1470 will be displayed only once.
1471 .RE

1473 .RE

1475 .sp
1476 .ne 2
1477 .na
1478 \fB\fBdladm create-bridge\fR [ \fB-P\fR \fIprioritize\fR] [ \fB-R\fR
1479 [ \fIroot-dir\fR] [ \fB-p\fR \fIprioritize\fR] [ \fB-m\fR \fImax-age\fR] [ \fB-h\fR
1480 [ \fIhello-time\fR] [ \fB-d\fR \fIfwd-delay\fR] [ \fB-f\fR
1481 [ \fIforce-protocol\fR] [ \fB-l\fR \fIlink\fR...] \fIbridge-name\fR\fR
1482 .ad
1483 .sp .6
1484 .RS 4n
1485 Create an 802.1D bridge instance and optionally assign one or more network
1486 links to the new bridge. By default, no bridge instances are present on the
1487 system.
1488 .sp
1489 In order to bridge between links, you must create at least one bridge instance.
1490 Each bridge instance is separate, and there is no forwarding connection between
1491 bridges.
1492 .sp
1493 .ne 2
1494 .na
1495 \fB\fB-P\fR \fIprioritize\fR, \fB--prioritize\fR=\fIprioritize\fR\fR
1496 .ad
1497 .sp .6
1498 .RS 4n
1499 Specifies a protection method. The defined protection methods are \fBstp\fR for
1500 the Spanning Tree Protocol and trill for \fBTRILL\fR, which is used on
1501 Rbridges. The default value is \fBstp\fR.
1502 .RE

1504 .sp
1505 .ne 2
1506 .na
1507 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
1508 .ad
1509 .sp .6
1510 .RS 4n
1511 See "Options," above.
1512 .RE

```

```

1514 .sp
1515 .ne 2
1516 .na
1517 \fB\fB-p\fR \fIprioritize\fR, \fB--prioritize\fR=\fIprioritize\fR\fR
1518 .ad
1519 .sp .6
1520 .RS 4n
1521 Specifies the Bridge Priority. This sets the IEEE STP priority value for
1522 determining the root bridge node in the network. The default value is
1523 \fB32768\fR. Valid values are \fB0\fR (highest priority) to \fB61440\fR (lowest
1524 priority), in increments of 4096.
1525 .sp
1526 If a value not evenly divisible by 4096 is used, the system silently rounds
1527 downward to the next lower value that is divisible by 4096.
1528 .RE

1530 .sp
1531 .ne 2
1532 .na
1533 \fB\fB-m\fR \fImax-age\fR, \fB--max-age\fR=\fImax-age\fR\fR
1534 .ad
1535 .sp .6
1536 .RS 4n
1537 Specifies the maximum age for configuration information in seconds. This sets
1538 the STP Bridge Max Age parameter. This value is used for all nodes in the
1539 network if this node is the root bridge. Bridge link information older than
1540 this time is discarded. It defaults to 20 seconds. Valid values are from 6 to
1541 40 seconds. See the \fB-d\fR \fIfwd-delay\fR parameter for additional
1542 constraints.
1543 .RE

1545 .sp
1546 .ne 2
1547 .na
1548 \fB\fB-h\fR \fIhello-time\fR, \fB--hello-time\fR=\fIhello-time\fR\fR
1549 .ad
1550 .sp .6
1551 .RS 4n
1552 Specifies the STP Bridge Hello Time parameter. When this node is the root node,
1553 it sends Configuration BPDU's at this interval throughout the network. The
1554 default value is 2 seconds. Valid values are from 1 to 10 seconds. See the
1555 \fB-d\fR \fIfwd-delay\fR parameter for additional constraints.
1556 .RE

1558 .sp
1559 .ne 2
1560 .na
1561 \fB\fB-d\fR \fIfwd-delay\fR, \fB--forward-delay\fR=\fIfwd-delay\fR\fR
1562 .ad
1563 .sp .6
1564 .RS 4n
1565 Specifies the STP Bridge Forward Delay parameter. When this node is the root
1566 node, then all bridges in the network use this timer to sequence the link
1567 states when a port is enabled. The default value is 15 seconds. Valid values
1568 are from 4 to 30 seconds.
1569 .sp
1570 Bridges must obey the following two constraints:
1571 .sp
1572 .in +2
1573 .nf
1574 2 * (\fIfwd-delay\fR - 1.0) >= \fImax-age\fR
1576 \fImax-age\fR >= 2 * (\fIhello-time\fR + 1.0)
1577 .fi
1578 .in -2

```

```

1579 .sp
1581 Any parameter setting that would violate those constraints is treated as an
1582 error and causes the command to fail with a diagnostic message. The message
1583 provides valid alternatives to the supplied values.
1584 .RE

1586 .sp
1587 .ne 2
1588 .na
1589 \fB\fB-f\fR \fIforce-protocol\fR,
1590 \fB--force-protocol\fR=\fIforce-protocol\fR\fR
1591 .ad
1592 .sp .6
1593 .RS 4n
1594 Specifies the MSTP forced maximum supported protocol. The default value is 3.
1595 Valid values are non-negative integers. The current implementation does not
1596 support RSTP or MSTP, so this currently has no effect. However, to prevent MSTP
1597 from being used in the future, the parameter may be set to \fB0\fR for STP only
1598 or \fB2\fR for STP and RSTP.
1599 .RE

1601 .sp
1602 .ne 2
1603 .na
1604 \fB\fB-l\fR \fIlink\fR, \fB--link\fR=\fIlink\fR\fR
1605 .ad
1606 .sp .6
1607 .RS 4n
1608 Specifies one or more links to add to the newly-created bridge. This is similar
1609 to creating the bridge and then adding one or more links, as with the
1610 \fBadd-bridge\fR subcommand. However, if any of the links cannot be added, the
1611 entire command fails, and the new bridge itself is not created. To add multiple
1612 links on the same command line, repeat this option for each link. You are
1613 permitted to create bridges without links. For more information about link
1614 assignments, see the \fBadd-bridge\fR subcommand.
1615 .RE

1617 Bridge creation and link assignment require the \fBPRIV_SYS_DL_CONFIG\fR
1618 privilege. Bridge creation might fail if the optional bridging feature is not
1619 installed on the system.
1620 .RE

1622 .sp
1623 .ne 2
1624 .na
1625 \fB\fBdladm modify-bridge\fR [ \fB-P\fR \fIprotect\fR] [ \fB-R\fR
1626 [ \fIroot-dir\fR] [ \fB-p\fR \fIpriority\fR] [ \fB-m\fR \fImax-age\fR] [ \fB-h\fR
1627 [ \fIhello-time\fR] [ \fB-d\fR \fIfoward-delay\fR] [ \fB-f\fR
1628 \fIforce-protocol\fR] [ \fB-l\fR \fIlink\fR...] \fIbridge-name\fR\fR
1629 .ad
1630 .sp .6
1631 .RS 4n
1632 Modify the operational parameters of an existing bridge. The options are the
1633 same as for the \fBcreate-bridge\fR subcommand, except that the \fB-l\fR option
1634 is not permitted. To add links to an existing bridge, use the \fBadd-bridge\fR
1635 subcommand.
1636 .sp
1637 Bridge parameter modification requires the \fBPRIV_SYS_DL_CONFIG\fR privilege.
1638 .RE

1640 .sp
1641 .ne 2
1642 .na
1643 \fB\fBdladm delete-bridge\fR [ \fB-R\fR \fIroot-dir\fR] \fIbridge-name\fR\fR
1644 .ad

```

```

1645 .sp .6
1646 .RS 4n
1647 Delete a bridge instance. The bridge being deleted must not have any attached
1648 links. Use the \fBremove-bridge\fR subcommand to deactivate links before
1649 deleting a bridge.
1650 .sp
1651 Bridge deletion requires the \fBPRIV_SYS_DL_CONFIG\fR privilege.
1652 .sp
1653 The \fB-R\fR (\fB--root-dir\fR) option is the same as for the
1654 \fBcreate-bridge\fR subcommand.
1655 .RE

1657 .sp
1658 .ne 2
1659 .na
1660 \fB\fBdladm add-bridge\fR [ \fB-R\fR \fIroot-dir\fR] \fB-l\fR \fIlink\fR
1661 [ \fB-l\fR \fIlink\fR...] \fIbridge-name\fR\fR
1662 .ad
1663 .sp .6
1664 .RS 4n
1665 Add one or more links to an existing bridge. If multiple links are specified,
1666 and adding any one of them results in an error, the command fails and no
1667 changes are made to the system.
1668 .sp
1669 Link addition to a bridge requires the \fBPRIV_SYS_DL_CONFIG\fR privilege.
1670 .sp
1671 A link may be a member of at most one bridge. An error occurs when you attempt
1672 to add a link that already belongs to another bridge. To move a link from one
1673 bridge instance to another, remove it from the current bridge before adding it
1674 to a new one.
1675 .sp
1676 The links assigned to a bridge must not also be VLANs, VNICs, or tunnels. Only
1677 physical Ethernet datalinks, aggregation datalinks, wireless links, and
1678 Ethernet stubs are permitted to be assigned to a bridge.
1679 .sp
1680 Links assigned to a bridge must all have the same MTU. This is checked when the
1681 link is assigned. The link is added to the bridge in a deactivated form if it
1682 is not the first link on the bridge and it has a differing MTU.
1683 .sp
1684 Note that systems using bridging should not set the \fBeeprom\fR(1M)
1685 \fBlocal-mac-address?\fR variable to false.
1686 .sp
1687 The options are the same as for the \fBcreate-bridge\fR subcommand.
1688 .RE

1690 .sp
1691 .ne 2
1692 .na
1693 \fB\fBdladm remove-bridge\fR [ \fB-R\fR \fIroot-dir\fR] \fB-l\fR \fIlink\fR
1694 [ \fB-l\fR \fIlink\fR...] \fIbridge-name\fR\fR
1695 .ad
1696 .sp .6
1697 .RS 4n
1698 Remove one or more links from a bridge instance. If multiple links are
1699 specified, and removing any one of them would result in an error, the command
1700 fails and none are removed.
1701 .sp
1702 Link removal from a bridge requires the \fBPRIV_SYS_DL_CONFIG\fR privilege.
1703 .sp
1704 The options are the same as for the \fBcreate-bridge\fR subcommand.
1705 .RE

1707 .sp
1708 .ne 2
1709 .na
1710 \fB\fBdladm show-bridge\fR [ \fB-flt\fR] [ \fB-s\fR [ \fB-i\fR \fIinterval\fR]]
```

```

1711 [[\fB-p\fR \fB-o\fR \fIfield\fR,...] [\fIbridge-name\fR]\fR
1712 .ad
1713 .sp .6
1714 .RS 4n
1715 Show the running status and configuration of bridges, their attached links,
1716 learned forwarding entries, and \fbTRILL\fR nickname databases. When showing
1717 overall bridge status and configuration, the bridge name can be omitted to show
1718 all bridges. The other forms require a specified bridge.
1719 .sp
1720 The show-bridge subcommand accepts the following options:
1721 .sp
1722 .ne 2
1723 .na
1724 \fB\fB-i\fR \fIinterval\fR, \fB--interval\fR=\fIinterval\fR\fR
1725 .ad
1726 .sp .6
1727 .RS 4n
1728 Used with the \fB-s\fR option to specify an interval, in seconds, at which
1729 statistics should be displayed. If this option is not specified, statistics
1730 will be displayed only once.
1731 .RE

1733 .sp
1734 .ne 2
1735 .na
1736 \fB\fB-s\fR, \fB--statistics\fR\fR
1737 .ad
1738 .sp .6
1739 .RS 4n
1740 Display statistics for the specified bridges or for a given bridge's attached
1741 links. This option cannot be used with the \fB-f\fR and \fB-t\fR options.
1742 .RE

1744 .sp
1745 .ne 2
1746 .na
1747 \fB\fB-p\fR, \fB--parseable\fR\fR
1748 .ad
1749 .sp .6
1750 .RS 4n
1751 Display using a stable machine-parsable format. See "Parsable Output Format,"
1752 below.
1753 .RE

1755 .sp
1756 .ne 2
1757 .na
1758 \fB\fB-o\fR \fIfield\fR[...], \fB--output\fR=\fIfield\fR[...]\fR
1759 .ad
1760 .sp .6
1761 .RS 4n
1762 A case-insensitive, comma-separated list of output fields to display. The field
1763 names are described below. The special value all displays all fields. Each set
1764 of fields has its own default set to display when \fB-o\fR is not specified.
1765 .RE

1767 By default, the \fbshow-bridge\fR subcommand shows bridge configuration. The
1768 following fields can be shown:
1769 .sp
1770 .ne 2
1771 .na
1772 \fB\fBBRIDGE\fR\fR
1773 .ad
1774 .sp .6
1775 .RS 4n
1776 The name of the bridge.

```

```

1777 .RE
1779 .sp
1780 .ne 2
1781 .na
1782 \fB\fBADDRESS\fR\fR
1783 .ad
1784 .sp .6
1785 .RS 4n
1786 The Bridge Unique Identifier value (MAC address).
1787 .RE

1789 .sp
1790 .ne 2
1791 .na
1792 \fB\fBPRIORITY\fR\fR
1793 .ad
1794 .sp .6
1795 .RS 4n
1796 Configured priority value; set by \fB-p\fR with \fbcreate-bridge\fR and
1797 \fbmodify-bridge\fR.
1798 .RE

1800 .sp
1801 .ne 2
1802 .na
1803 \fB\fBMAXAGE\fR\fR
1804 .ad
1805 .sp .6
1806 .RS 4n
1807 Configured bridge maximum age; set by \fB-m\fR with \fbcreate-bridge\fR and
1808 \fbmodify-bridge\fR.
1809 .RE

1811 .sp
1812 .ne 2
1813 .na
1814 \fB\fBHELLOTIME\fR\fR
1815 .ad
1816 .sp .6
1817 .RS 4n
1818 Configured bridge hello time; set by \fB-h\fR with \fbcreate-bridge\fR and
1819 \fbmodify-bridge\fR.
1820 .RE

1822 .sp
1823 .ne 2
1824 .na
1825 \fB\fBFWDDELAY\fR\fR
1826 .ad
1827 .sp .6
1828 .RS 4n
1829 Configured forwarding delay; set by \fB-d\fR with \fbcreate-bridge\fR and
1830 \fbmodify-bridge\fR.
1831 .RE

1833 .sp
1834 .ne 2
1835 .na
1836 \fB\fBFORCEPROTO\fR\fR
1837 .ad
1838 .sp .6
1839 .RS 4n
1840 Configured forced maximum protocol; set by \fB-f\fR with \fbcreate-bridge\fR
1841 and \fbmodify-bridge\fR.
1842 .RE

```

```

1844 .sp
1845 .ne 2
1846 .na
1847 \fB\fBTCTIME\fR\fR
1848 .ad
1849 .sp .6
1850 .RS 4n
1851 Time, in seconds, since last topology change.
1852 .RE

1854 .sp
1855 .ne 2
1856 .na
1857 \fB\fBTCCOUNT\fR\fR
1858 .ad
1859 .sp .6
1860 .RS 4n
1861 Count of the number of topology changes.
1862 .RE

1864 .sp
1865 .ne 2
1866 .na
1867 \fB\fBTCHANGE\fR\fR
1868 .ad
1869 .sp .6
1870 .RS 4n
1871 This indicates that a topology change was detected.
1872 .RE

1874 .sp
1875 .ne 2
1876 .na
1877 \fB\fBDESROOT\fR\fR
1878 .ad
1879 .sp .6
1880 .RS 4n
1881 Bridge Identifier of the root node.
1882 .RE

1884 .sp
1885 .ne 2
1886 .na
1887 \fB\fBROOTCOST\fR\fR
1888 .ad
1889 .sp .6
1890 .RS 4n
1891 Cost of the path to the root node.
1892 .RE

1894 .sp
1895 .ne 2
1896 .na
1897 \fB\fBROOTPORT\fR\fR
1898 .ad
1899 .sp .6
1900 .RS 4n
1901 Port number used to reach the root node.
1902 .RE

1904 .sp
1905 .ne 2
1906 .na
1907 \fB\fBMAXAGE\fR\fR
1908 .ad

```

```

1909 .sp .6
1910 .RS 4n
1911 Maximum age value from the root node.
1912 .RE

1914 .sp
1915 .ne 2
1916 .na
1917 \fB\fBHELLOTIME\fR\fR
1918 .ad
1919 .sp .6
1920 .RS 4n
1921 Hello time value from the root node.
1922 .RE

1924 .sp
1925 .ne 2
1926 .na
1927 \fB\fBFWDDELAY\fR\fR
1928 .ad
1929 .sp .6
1930 .RS 4n
1931 Forward delay value from the root node.
1932 .RE

1934 .sp
1935 .ne 2
1936 .na
1937 \fB\fBHOLDTIME\fR\fR
1938 .ad
1939 .sp .6
1940 .RS 4n
1941 Minimum BPDU interval.
1942 .RE

1944 By default, when the \fB-o\fR option is not specified, only the \fBBRIDGE\fR,
1945 \fBADDRESS\fR, \fBPRIORITY\fR, and \fBDESROOT\fR fields are shown.
1946 .sp
1947 When the \fB-s\fR option is specified, the \fBshow-bridge\fR subcommand shows
1948 bridge statistics. The following fields can be shown:
1949 .sp
1950 .ne 2
1951 .na
1952 \fB\fBBRIDGE\fR\fR
1953 .ad
1954 .sp .6
1955 .RS 4n
1956 Bridge name.
1957 .RE

1959 .sp
1960 .ne 2
1961 .na
1962 \fB\fBDROPS\fR\fR
1963 .ad
1964 .sp .6
1965 .RS 4n
1966 Number of packets dropped due to resource problems.
1967 .RE

1969 .sp
1970 .ne 2
1971 .na
1972 \fB\fBFORWARDS\fR\fR
1973 .ad
1974 .sp .6

```

```

1975 .RS 4n
1976 Number of packets forwarded from one link to another.
1977 .RE

1979 .sp
1980 .ne 2
1981 .na
1982 \fB\fBMBCAST\fR\fR
1983 .ad
1984 .sp .6
1985 .RS 4n
1986 Number of multicast and broadcast packets handled by the bridge.
1987 .RE

```

```

1989 .sp
1990 .ne 2
1991 .na
1992 \fB\fBRECV\fR\fR
1993 .ad
1994 .sp .6
1995 .RS 4n
1996 Number of packets received on all attached links.
1997 .RE

```

```

1999 .sp
2000 .ne 2
2001 .na
2002 \fB\fBSENT\fR\fR
2003 .ad
2004 .sp .6
2005 .RS 4n
2006 Number of packets sent on all attached links.
2007 .RE

```

```

2009 .sp
2010 .ne 2
2011 .na
2012 \fB\fBUKNOWN\fR\fR
2013 .ad
2014 .sp .6
2015 .RS 4n
2016 Number of packets handled that have an unknown destination. Such packets are
2017 sent to all links.
2018 .RE

```

```

2020 By default, when the \fB-o\fR option is not specified, only the \fBBRIDGE\fR,
2021 \fBDROPS\fR, and \fBFORWARDS\fR fields are shown.
2022 .sp
2023 The \fBshow-bridge\fR subcommand also accepts the following options:
2024 .sp
2025 .ne 2
2026 .na
2027 \fB\fB-1\fR, \fB--link\fR\fR
2028 .ad
2029 .sp .6
2030 .RS 4n
2031 Displays link-related status and statistics information for all links attached
2032 to a single bridge instance. By using this option and without the \fB-s\fR
2033 option, the following fields can be displayed for each link:
2034 .sp
2035 .ne 2
2036 .na
2037 \fB\fBLINK\fR\fR
2038 .ad
2039 .sp .6
2040 .RS 4n

```

```

2041 The link name.
2042 .RE

2044 .sp
2045 .ne 2
2046 .na
2047 \fB\fBINDEX\fR\fR
2048 .ad
2049 .sp .6
2050 .RS 4n
2051 Port (link) index number on the bridge.
2052 .RE

2054 .sp
2055 .ne 2
2056 .na
2057 \fB\fBSTATE\fR\fR
2058 .ad
2059 .sp .6
2060 .RS 4n
2061 State of the link. The state can be \fBdisabled\fR, \fBdiscarding\fR,
2062 \fBlearning\fR, \fBforwarding\fR, \fBnon-stp\fR, or \fBbad-mtu\fR.
2063 .RE

2065 .sp
2066 .ne 2
2067 .na
2068 \fB\fBUPTIME\fR\fR
2069 .ad
2070 .sp .6
2071 .RS 4n
2072 Number of seconds since the last reset or initialization.
2073 .RE

2075 .sp
2076 .ne 2
2077 .na
2078 \fB\fBOPER COST\fR\fR
2079 .ad
2080 .sp .6
2081 .RS 4n
2082 Actual cost in use (1-65535).
2083 .RE

2085 .sp
2086 .ne 2
2087 .na
2088 \fB\fBOPER P2P\fR\fR
2089 .ad
2090 .sp .6
2091 .RS 4n
2092 This indicates whether point-to-point (\fBP2P\fR) mode been detected.
2093 .RE

2095 .sp
2096 .ne 2
2097 .na
2098 \fB\fBOPER EDGE\fR\fR
2099 .ad
2100 .sp .6
2101 .RS 4n
2102 This indicates whether edge mode has been detected.
2103 .RE

2105 .sp
2106 .ne 2

```

```

2107 .na
2108 \fB\fBDESROOT\fR\fR
2109 .ad
2110 .sp .6
2111 .RS 4n
2112 The Root Bridge Identifier that has been seen on this port.
2113 .RE

2115 .sp
2116 .ne 2
2117 .na
2118 \fB\fBDESCOST\fR\fR
2119 .ad
2120 .sp .6
2121 .RS 4n
2122 Path cost to the network root node through the designated port.
2123 .RE

2125 .sp
2126 .ne 2
2127 .na
2128 \fB\fBDESBRIDGE\fR\fR
2129 .ad
2130 .sp .6
2131 .RS 4n
2132 Bridge Identifier for this port.
2133 .RE

2135 .sp
2136 .ne 2
2137 .na
2138 \fB\fBDESPORT\fR\fR
2139 .ad
2140 .sp .6
2141 .RS 4n
2142 The ID and priority of the port used to transmit configuration messages for
2143 this port.
2144 .RE

2146 .sp
2147 .ne 2
2148 .na
2149 \fB\fBTCACK\fR\fR
2150 .ad
2151 .sp .6
2152 .RS 4n
2153 This indicates whether Topology Change Acknowledge has been seen.
2154 .RE

2156 When the \fB-l\fR option is specified without the \fB-o\fR option, only the
2157 \fBLINK\fR, \fBSTATE\fR, \fBUPTIME\fR, and \fBDESROOT\fR fields are shown.
2158 .sp
2159 When the \fB-l\fR option is specified, the \fB-s\fR option can be used to
2160 display the following fields for each link:
2161 .sp
2162 .ne 2
2163 .na
2164 \fB\fBLINK\fR\fR
2165 .ad
2166 .sp .6
2167 .RS 4n
2168 Link name.
2169 .RE

2171 .sp
2172 .ne 2

```

```

2173 .na
2174 \fB\fBCFGBPDU\fR\fR
2175 .ad
2176 .sp .6
2177 .RS 4n
2178 Number of configuration BPDUs received.
2179 .RE

2181 .sp
2182 .ne 2
2183 .na
2184 \fB\fBTCBPDU\fR\fR
2185 .ad
2186 .sp .6
2187 .RS 4n
2188 Number of topology change BPDUs received.
2189 .RE

2191 .sp
2192 .ne 2
2193 .na
2194 \fB\fBRSTPBPDU\fR\fR
2195 .ad
2196 .sp .6
2197 .RS 4n
2198 Number of Rapid Spanning Tree BPDUs received.
2199 .RE

2201 .sp
2202 .ne 2
2203 .na
2204 \fB\fBTXPDU\fR\fR
2205 .ad
2206 .sp .6
2207 .RS 4n
2208 Number of BPDUs transmitted.
2209 .RE

2211 .sp
2212 .ne 2
2213 .na
2214 \fB\fBDROPS\fR\fR
2215 .ad
2216 .sp .6
2217 .RS 4n
2218 Number of packets dropped due to resource problems.
2219 .RE

2221 .sp
2222 .ne 2
2223 .na
2224 \fB\fBRECV\fR\fR
2225 .ad
2226 .sp .6
2227 .RS 4n
2228 Number of packets received by the bridge.
2229 .RE

2231 .sp
2232 .ne 2
2233 .na
2234 \fB\fBXMIT\fR\fR
2235 .ad
2236 .sp .6
2237 .RS 4n
2238 Number of packets sent by the bridge.

```

```

2239 .RE
2241 When the \fB-o\fR option is not specified, only the \fBLINK\fR, \fBDROPS\fR,
2242 \fBRECV\fR, and \fBXMIT\fR fields are shown.
2243 .RE

2245 .sp
2246 .ne 2
2247 .na
2248 \fB\fB-f\fR, \fB--forwarding\fR\fR
2249 .ad
2250 .sp .6
2251 .RS 4n
2252 Displays forwarding entries for a single bridge instance. With this option, the
2253 following fields can be shown for each forwarding entry:
2254 .sp
2255 .ne 2
2256 .na
2257 \fB\fBDEST\fR\fR
2258 .ad
2259 .sp .6
2260 .RS 4n
2261 Destination MAC address.
2262 .RE

2264 .sp
2265 .ne 2
2266 .na
2267 \fB\fBAGE\fR\fR
2268 .ad
2269 .sp .6
2270 .RS 4n
2271 Age of entry in seconds and milliseconds. Omitted for local entries.
2272 .RE

2274 .sp
2275 .ne 2
2276 .na
2277 \fB\fBFLAGS\fR\fR
2278 .ad
2279 .sp .6
2280 .RS 4n
2281 The \fBL\fR (local) flag is shown if the MAC address belongs to an attached
2282 link or to a VNIC on one of the attached links.
2283 .RE

2285 .sp
2286 .ne 2
2287 .na
2288 \fB\fBOUTPUT\fR\fR
2289 .ad
2290 .sp .6
2291 .RS 4n
2292 For local entries, this is the name of the attached link that has the MAC
2293 address. Otherwise, for bridges that use Spanning Tree Protocol, this is the
2294 output interface name. For R Bridges, this is the output \fBTILL\fR nickname.
2295 .RE

2297 When the \fB-o\fR option is not specified, the \fBDEST\fR, \fBAGE\fR,
2298 \fBFLAGS\fR, and \fBOUTPUT\fR fields are shown.
2299 .RE

2301 .sp
2302 .ne 2
2303 .na
2304 \fB\fB-t\fR, \fB--trill\fR\fR

```

```

2305 .ad
2306 .sp .6
2307 .RS 4n
2308 Displays \fBTILL\fR nickname entries for a single bridge instance. With this
2309 option, the following fields can be shown for each \fBTILL\fR nickname entry:
2310 .sp
2311 .ne 2
2312 .na
2313 \fB\fBNICK\fR\fR
2314 .ad
2315 .sp .6
2316 .RS 4n
2317 \fBTILL\fR nickname for this RBridge, which is a number from 1 to 65535.
2318 .RE

2320 .sp
2321 .ne 2
2322 .na
2323 \fB\fBFLAGS\fR\fR
2324 .ad
2325 .sp .6
2326 .RS 4n
2327 The \fBL\fR flag is shown if the nickname identifies the local system.
2328 .RE

2330 .sp
2331 .ne 2
2332 .na
2333 \fB\fBLINK\fR\fR
2334 .ad
2335 .sp .6
2336 .RS 4n
2337 Link name for output when sending messages to this RBridge.
2338 .RE

2340 .sp
2341 .ne 2
2342 .na
2343 \fB\fBNEXTHOP\fR\fR
2344 .ad
2345 .sp .6
2346 .RS 4n
2347 MAC address of the next hop RBridge that is used to reach the RBridge with this
2348 nickname.
2349 .RE

2351 When the \fB-o\fR option is not specified, the \fBNICK\fR, \fBFLAGS\fR,
2352 \fBLINK\fR, and \fBNEXTHOP\fR fields are shown.
2353 .RE

2355 .RE

2357 .sp
2358 .ne 2
2359 .na
2360 \fB\fBdadm create-vlan\fR [\fB-ft\fR] [\fB-R\fR \fIroot-dir\fR] \fB-l\fR
2361 \fIether-link\fR \fB-v\fR \fIvid\fR [\fIvlan-link\fR]\fR
2362 .ad
2363 .sp .6
2364 .RS 4n
2365 Create a tagged VLAN link with an ID of \fIvid\fR over Ethernet link
2366 \fIether-link\fR. The name of the VLAN link can be specified as
2367 \fIvlan\fR-\fIlink\fR. If the name is not specified, a name will be
2368 automatically generated (assuming that \fIether-link\fR is \fIname\fR\fIPPA\fR)
2369 as:
2370 .sp

```

```

2371 .in +2
2372 .nf
2373 <\fIname\fR><1000 * \fIvlan-tag\fR + \fIPPA\fR>
2374 .fi
2375 .in -2
2376 .sp

2378 For example, if \fIether-link\fR is \fBbge1\fR and \fIvid\fR is 2, the name
2379 generated is \fBbge2001\fR.
2380 .sp
2381 .ne 2
2382 .na
2383 \fB\fB-f\fR, \fB--force\fR\fR
2384 .ad
2385 .sp .6
2386 .RS 4n
2387 Force the creation of the VLAN link. Some devices do not allow frame sizes
2388 large enough to include a VLAN header. When creating a VLAN link over such a
2389 device, the \fB-f\fR option is needed, and the MTU of the IP interfaces on the
2390 resulting VLAN must be set to 1496 instead of 1500.
2391 .RE

2393 .sp
2394 .ne 2
2395 .na
2396 \fB\fB-l\fR \fIether-link\fR\fR
2397 .ad
2398 .sp .6
2399 .RS 4n
2400 Specifies Ethernet link over which VLAN is created.
2401 .RE

2403 .sp
2404 .ne 2
2405 .na
2406 \fB\fB-t\fR, \fB--temporary\fR\fR
2407 .ad
2408 .sp .6
2409 .RS 4n
2410 Specifies that the VLAN link is temporary. Temporary VLAN links last until the
2411 next reboot.
2412 .RE

2414 .sp
2415 .ne 2
2416 .na
2417 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
2418 .ad
2419 .sp .6
2420 .RS 4n
2421 See "Options," above.
2422 .RE

2424 .RE

2426 .sp
2427 .ne 2
2428 .na
2429 \fB\fBdladm delete-vlan\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR]
2430 \fIvlan-link\fR\fR
2431 .ad
2432 .sp .6
2433 .RS 4n
2434 Delete the VLAN link specified.
2435 .sp
2436 The \fBdelete-vlan\fRsubcommand accepts the following options:

```

```

2437 .sp
2438 .ne 2
2439 .na
2440 \fB\fB-t\fR, \fB--temporary\fR\fR
2441 .ad
2442 .sp .6
2443 .RS 4n
2444 Specifies that the deletion is temporary. Temporary deletions last until the
2445 next reboot.
2446 .RE

2448 .sp
2449 .ne 2
2450 .na
2451 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
2452 .ad
2453 .sp .6
2454 .RS 4n
2455 See "Options," above.
2456 .RE

2458 .RE

2460 .sp
2461 .ne 2
2462 .na
2463 \fB\fBdladm show-vlan\fR [\fB-P\fR] [[\fB-p\fR] \fB-o\fR \fIfield\fR[,...]]
2464 [\fIvlan-link\fR]\fR
2465 .ad
2466 .sp .6
2467 .RS 4n
2468 Display VLAN configuration for all VLAN links or for the specified VLAN link.
2469 .sp
2470 The \fBshow-vlan\fRsubcommand accepts the following options:
2471 .sp
2472 .ne 2
2473 .na
2474 \fB\fB-o\fR \fIfield\fR[,...], \fB--output\fR=\fIfield\fR[,...]\fR
2475 .ad
2476 .sp .6
2477 .RS 4n
2478 A case-insensitive, comma-separated list of output fields to display. The field
2479 name must be one of the fields listed below, or the special value \fBall\fR, to
2480 display all fields. For each VLAN link, the following fields can be displayed:
2481 .sp
2482 .ne 2
2483 .na
2484 \fB\fBLINK\fR\fR
2485 .ad
2486 .sp .6
2487 .RS 4n
2488 The name of the VLAN link.
2489 .RE

2491 .sp
2492 .ne 2
2493 .na
2494 \fB\fBVID\fR\fR
2495 .ad
2496 .sp .6
2497 .RS 4n
2498 The ID associated with the VLAN.
2499 .RE

2501 .sp
2502 .ne 2

```

```

2503 .na
2504 \fB\fBOVER\fR\fR
2505 .ad
2506 .sp .6
2507 .RS 4n
2508 The name of the physical link over which this VLAN is configured.
2509 .RE

2511 .sp
2512 .ne 2
2513 .na
2514 \fB\fBFLAGS\fR\fR
2515 .ad
2516 .sp .6
2517 .RS 4n
2518 A set of flags associated with the VLAN link. Possible flags are:
2519 .sp
2520 .ne 2
2521 .na
2522 \fB\fBf\fR\fR
2523 .ad
2524 .sp .6
2525 .RS 4n
2526 The VLAN was created using the \fB-f\fR option to \fBcreate-vlan\fR.
2527 .RE

2529 .sp
2530 .ne 2
2531 .na
2532 \fB\fBi\fR\fR
2533 .ad
2534 .sp .6
2535 .RS 4n
2536 The VLAN was implicitly created when the DLPI link was opened. These VLAN links
2537 are automatically deleted on last close of the DLPI link (for example, when the
2538 IP interface associated with the VLAN link is unplumbed).
2539 .RE

2541 Additional flags might be defined in the future.
2542 .RE

2544 .RE

2546 .sp
2547 .ne 2
2548 .na
2549 \fB\fB-p\fR, \fB--parseable\fR\fR
2550 .ad
2551 .sp .6
2552 .RS 4n
2553 Display using a stable machine-parseable format. The \fB-o\fR option is
2554 required with \fB-p\fR. See "Parseable Output Format", below.
2555 .RE

2557 .sp
2558 .ne 2
2559 .na
2560 \fB\fB-P\fR, \fB--persistent\fR\fR
2561 .ad
2562 .sp .6
2563 .RS 4n
2564 Display the persistent VLAN configuration rather than the state of the running
2565 system.
2566 .RE
2568 .RE

```

```

2570 .sp
2571 .ne 2
2572 .na
2573 \fB\fBdladm scan-wifi\fR [[\fB-p\fR] \fB-o\fR \fIifield\fR[,...]]
2574 [\fIwifi-link\fR]\fR
2575 .ad
2576 .sp .6
2577 .RS 4n
2578 Scans for \fBWiFi\fR networks, either on all \fBWiFi\fR links, or just on the
2579 specified \fIwifi-link\fR.
2580 .sp
2581 By default, currently all fields but \fBSSTYPE\fR are displayed.
2582 .sp
2583 .ne 2
2584 .na
2585 \fB\fB-o\fR \fIifield\fR[,...], \fB--output\fR=\fIifield\fR[,...]\fR
2586 .ad
2587 .sp .6
2588 .RS 4n
2589 A case-insensitive, comma-separated list of output fields to display. The field
2590 name must be one of the fields listed below, or the special value \fBall\fR to
2591 display all fields. For each \fBWiFi\fR network found, the following fields can
2592 be displayed:
2593 .sp
2594 .ne 2
2595 .na
2596 \fB\fBLINK\fR\fR
2597 .ad
2598 .sp .6
2599 .RS 4n
2600 The name of the link the \fBWiFi\fR network is on.
2601 .RE

2603 .sp
2604 .ne 2
2605 .na
2606 \fB\fBESSID\fR\fR
2607 .ad
2608 .sp .6
2609 .RS 4n
2610 The \fBESSID\fR (name) of the \fBWiFi\fR network.
2611 .RE

2613 .sp
2614 .ne 2
2615 .na
2616 \fB\fBBSSID\fR\fR
2617 .ad
2618 .sp .6
2619 .RS 4n
2620 Either the hardware address of the \fBWiFi\fR network's Access Point (for
2621 \fBBSS\fR networks), or the \fBWiFi\fR network's randomly generated unique
2622 token (for \fBIBSS\fR networks).
2623 .RE

2625 .sp
2626 .ne 2
2627 .na
2628 \fB\fBSEC\fR\fR
2629 .ad
2630 .sp .6
2631 .RS 4n
2632 Either \fBnone\fR for a \fBWiFi\fR network that uses no security, \fBwep\fR for
2633 a \fBWiFi\fR network that requires WEP (Wired Equivalent Privacy), or \fBwpa\fR
2634 for a WiFi network that requires WPA (Wi-Fi Protected Access).

```

```

2635 .RE
2637 .sp
2638 .ne 2
2639 .na
2640 \fB\fBMODE\fR\fR
2641 .ad
2642 .sp .6
2643 .RS 4n
2644 The supported connection modes: one or more of \fBa\fR, \fBb\fR, or \fBg\fR.
2645 .RE

2647 .sp
2648 .ne 2
2649 .na
2650 \fB\fBSTRENGTH\fR\fR
2651 .ad
2652 .sp .6
2653 .RS 4n
2654 The strength of the signal: one of \fBexcellent\fR, \fBvery good\fR,
2655 \fBgood\fR, \fBweak\fR, or \fBvery weak\fR.
2656 .RE

2658 .sp
2659 .ne 2
2660 .na
2661 \fB\fBSPEED\fR\fR
2662 .ad
2663 .sp .6
2664 .RS 4n
2665 The maximum speed of the \fBWiFi\fR network, in megabits per second.
2666 .RE

2668 .sp
2669 .ne 2
2670 .na
2671 \fB\fBSSTYPE\fR\fR
2672 .ad
2673 .sp .6
2674 .RS 4n
2675 Either \fBbss\fR for \fBBSS\fR (infrastructure) networks, or \fBibss\fR for
2676 \fBIBSS\fR (ad-hoc) networks.
2677 .RE

2679 .RE

2681 .sp
2682 .ne 2
2683 .na
2684 \fB\fB-p\fR, \fB--parseable\fR\fR
2685 .ad
2686 .sp .6
2687 .RS 4n
2688 Display using a stable machine-parseable format. The \fB-o\fR option is
2689 required with \fB-p\fR. See "Parseable Output Format", below.
2690 .RE

2692 .RE

2694 .sp
2695 .ne 2
2696 .na
2697 \fB\fBdladm connect-wifi\fR [\fB-e\fR \fIessid\fR] [\fB-i\fR \fIbssid\fR]
2698 [\fB-k\fR \fIkey\fR,...] [\fB-s\fR \fBnone\fR | \fBwep\fR | \fBwpa\fR]
2699 [\fB-a\fR \fBopen\fR|\fBshared\fR] [\fB-b\fR \fBbss\fR|\fBibss\fR] [\fB-c\fR]
2700 [\fB-m\fR \fBa\fR|\fBb\fR|\fBg\fR] [\fB-T\fR \fItime\fR] [\fIwifi-link\fR]\fR

```

```

2701 .ad
2702 .sp .6
2703 .RS 4n
2704 Connects to a \fBWiFi\fR network. This consists of four steps: \fIdiscovery\fR,
2705 \fIfiltration\fR, \fIprioritization\fR, and \fIassociation\fR. However, to
2706 enable connections to non-broadcast \fBWiFi\fR networks and to improve
2707 performance, if a \fBBSSID\fR or \fBESSID\fR is specified using the \fB-e\fR or
2708 \fB-i\fR options, then the first three steps are skipped and \fBconnect-wifi\fR
2709 immediately attempts to associate with a \fBBSSID\fR or \fBESSID\fR that
2710 matches the rest of the provided parameters. If this association fails, but
2711 there is a possibility that other networks matching the specified criteria
2712 exist, then the traditional discovery process begins as specified below.
2713 .sp
2714 The discovery step finds all available \fBWiFi\fR networks on the specified
2715 WiFi link, which must not yet be connected. For administrative convenience, if
2716 there is only one \fBWiFi\fR link on the system, \fIwifi-link\fR can be
2717 omitted.
2718 .sp
2719 Once discovery is complete, the list of networks is filtered according to the
2720 value of the following options:
2721 .sp
2722 .ne 2
2723 .na
2724 \fB\fB-e\fR \fIessid\fR \fB--essid\fR=\fIessid\fR\fR
2725 .ad
2726 .sp .6
2727 .RS 4n
2728 Networks that do not have the same \fIessid\fR are filtered out.
2729 .RE

2731 .sp
2732 .ne 2
2733 .na
2734 \fB\fB-b\fR \fBbss\fR|\fBibss\fR, \fB--bsstype\fR=\fBbss\fR|\fBibss\fR\fR
2735 .ad
2736 .sp .6
2737 .RS 4n
2738 Networks that do not have the same \fBbsstype\fR are filtered out.
2739 .RE

2741 .sp
2742 .ne 2
2743 .na
2744 \fB\fB-m\fR \fBa\fR|\fBb\fR|\fBg\fR, \fB--mode\fR=\fBa\fR|\fBb\fR|\fBg\fR\fR
2745 .ad
2746 .sp .6
2747 .RS 4n
2748 Networks not appropriate for the specified 802.11 mode are filtered out.
2749 .RE

2751 .sp
2752 .ne 2
2753 .na
2754 \fB\fB-k\fR \fIkey\fR,...\fR, \fB--key\fR=\fIkey\fR,...\fR\fR
2755 .ad
2756 .sp .6
2757 .RS 4n
2758 Use the specified \fBsecobj\fR named by the key to connect to the network.
2759 Networks not appropriate for the specified keys are filtered out.
2760 .RE

2762 .sp
2763 .ne 2
2764 .na
2765 \fB\fB-s\fR \fBnone\fR|\fBwep\fR|\fBwpa\fR,
2766 \fB--sec\fR=\fBnone\fR|\fBwep\fR|\fBwpa\fR\fR

```

```

2767 .ad
2768 .sp .6
2769 .RS 4n
2770 Networks not appropriate for the specified security mode are filtered out.
2771 .RE

2773 Next, the remaining networks are prioritized, first by signal strength, and
2774 then by maximum speed. Finally, an attempt is made to associate with each
2775 network in the list, in order, until one succeeds or no networks remain.
2776 .sp
2777 In addition to the options described above, the following options also control
2778 the behavior of \fBconnect-wifi\fR:
2779 .sp
2780 .ne 2
2781 .na
2782 \fB\fB-a\fR \fBopen\fR|\fBshared\fR, \fB--auth\fR=\fBopen\fR|\fBshared\fR\fR
2783 .ad
2784 .sp .6
2785 .RS 4n
2786 Connect using the specified authentication mode. By default, \fBopen\fR and
2787 \fBshared\fR are tried in order.
2788 .RE

2790 .sp
2791 .ne 2
2792 .na
2793 \fB\fB-c\fR, \fB--create-ibss\fR\fR
2794 .ad
2795 .sp .6
2796 .RS 4n
2797 Used with \fB-b ibss\fR to create a new ad-hoc network if one matching the
2798 specified \fBESSID\fR cannot be found. If no \fBESSID\fR is specified, then
2799 \fB-c -b ibss\fR always triggers the creation of a new ad-hoc network.
2800 .RE

2802 .sp
2803 .ne 2
2804 .na
2805 \fB\fB-T\fR \fIitime\fR, \fB--timeout\fR=\fIitime\fR\fR
2806 .ad
2807 .sp .6
2808 .RS 4n
2809 Specifies the number of seconds to wait for association to succeed. If
2810 \fIitime\fR is \fBforever\fR, then the associate will wait indefinitely. The
2811 current default is ten seconds, but this might change in the future. Timeouts
2812 shorter than the default might not succeed reliably.
2813 .RE

2815 .sp
2816 .ne 2
2817 .na
2818 \fB\fB-k\fR \fIkey,...\fR, \fB--key\fR=\fIkey,...\fR\fR
2819 .ad
2820 .sp .6
2821 .RS 4n
2822 In addition to the filtering previously described, the specified keys will be
2823 used to secure the association. The security mode to use will be based on the
2824 key class; if a security mode was explicitly specified, it must be compatible
2825 with the key class. All keys must be of the same class.
2826 .sp
2827 For security modes that support multiple key slots, the slot to place the key
2828 will be specified by a colon followed by an index. Therefore, \fB-k mykey:3\fR
2829 places \fBmykey\fR in slot 3. By default, slot 1 is assumed. For security modes
2830 that support multiple keys, a comma-separated list can be specified, with the
2831 first key being the active key.
2832 .RE

```

```

2834 .RE
2836 .sp
2837 .ne 2
2838 .na
2839 \fB\fBdladm disconnect-wifi\fR [\fB-a\fR] [\fIwifi-link\fR]\fR
2840 .ad
2841 .sp .6
2842 .RS 4n
2843 Disconnect from one or more \fBWIFI\fR networks. If \fIwifi-link\fR specifies a
2844 connected \fBWIFI\fR link, then it is disconnected. For administrative
2845 convenience, if only one \fBWIFI\fR link is connected, \fIwifi-link\fR can be
2846 omitted.
2847 .sp
2848 .ne 2
2849 .na
2850 \fB\fB-a\fR, \fB--all-links\fR\fR
2851 .ad
2852 .sp .6
2853 .RS 4n
2854 Disconnects from all connected links. This is primarily intended for use by
2855 scripts.
2856 .RE

2858 .RE
2860 .sp
2861 .ne 2
2862 .na
2863 \fB\fBdladm show-wifi\fR [[\fB-p\fR] \fB-o\fR \fIfield\fR,...]
2864 [\fIwifi-link\fR]\fR
2865 .ad
2866 .sp .6
2867 .RS 4n
2868 Shows \fBWIFI\fR configuration information either for all \fBWIFI\fR links or
2869 for the specified link \fIwifi-link\fR.
2870 .sp
2871 .ne 2
2872 .na
2873 \fB\fB-o\fR \fIfield,...\fR, \fB--output\fR=\fIfield\fR\fR
2874 .ad
2875 .sp .6
2876 .RS 4n
2877 A case-insensitive, comma-separated list of output fields to display. The field
2878 name must be one of the fields listed below, or the special value \fBall\fR, to
2879 display all fields. For each \fBWIFI\fR link, the following fields can be
2880 displayed:
2881 .sp
2882 .ne 2
2883 .na
2884 \fB\fBLINK\fR\fR
2885 .ad
2886 .sp .6
2887 .RS 4n
2888 The name of the link being displayed.
2889 .RE

2891 .sp
2892 .ne 2
2893 .na
2894 \fB\fBSTATUS\fR\fR
2895 .ad
2896 .sp .6
2897 .RS 4n
2898 Either \fBconnected\fR if the link is connected, or \fBdisconnected\fR if it is

```

2899 not connected. If the link is disconnected, all remaining fields have the value  
 2900 \fB--\fR.  
 2901 .RE

2903 .sp  
 2904 .ne 2  
 2905 .na  
 2906 \fB\fBESSID\fR\fR  
 2907 .ad  
 2908 .sp .6  
 2909 .RS 4n  
 2910 The \fBESSID\fR (name) of the connected \fBWiFi\fR network.  
 2911 .RE

2913 .sp  
 2914 .ne 2  
 2915 .na  
 2916 \fB\fBBSSID\fR\fR  
 2917 .ad  
 2918 .sp .6  
 2919 .RS 4n  
 2920 Either the hardware address of the \fBWiFi\fR network's Access Point (for  
 2921 \fBBSS\fR networks), or the \fBWiFi\fR network's randomly generated unique  
 2922 token (for \fBIBSS\fR networks).  
 2923 .RE

2925 .sp  
 2926 .ne 2  
 2927 .na  
 2928 \fB\fBSEC\fR\fR  
 2929 .ad  
 2930 .sp .6  
 2931 .RS 4n  
 2932 Either \fBnone\fR for a \fBWiFi\fR network that uses no security, \fBwep\fR for  
 2933 a \fBWiFi\fR network that requires WEP, or \fBwpa\fR for a WiFi network that  
 2934 requires WPA.  
 2935 .RE

2937 .sp  
 2938 .ne 2  
 2939 .na  
 2940 \fB\fBMODE\fR\fR  
 2941 .ad  
 2942 .sp .6  
 2943 .RS 4n  
 2944 The supported connection modes: one or more of \fBa\fR, \fBb\fR, or \fBg\fR.  
 2945 .RE

2947 .sp  
 2948 .ne 2  
 2949 .na  
 2950 \fB\fBSTRENGTH\fR\fR  
 2951 .ad  
 2952 .sp .6  
 2953 .RS 4n  
 2954 The connection strength: one of \fBexcellent\fR, \fBvery good\fR, \fBgood\fR,  
 2955 \fBweak\fR, or \fBvery weak\fR.  
 2956 .RE

2958 .sp  
 2959 .ne 2  
 2960 .na  
 2961 \fB\fBSPEED\fR\fR  
 2962 .ad  
 2963 .sp .6  
 2964 .RS 4n

2965 The connection speed, in megabits per second.  
 2966 .RE

2968 .sp  
 2969 .ne 2  
 2970 .na  
 2971 \fB\fBAUTH\fR\fR  
 2972 .ad  
 2973 .sp .6  
 2974 .RS 4n  
 2975 Either \fBopen\fR or \fBshared\fR (see \fBconnect-wifi\fR).  
 2976 .RE

2978 .sp  
 2979 .ne 2  
 2980 .na  
 2981 \fB\fBSSSTYPE\fR\fR  
 2982 .ad  
 2983 .sp .6  
 2984 .RS 4n  
 2985 Either \fBbss\fR for \fBBSS\fR (infrastructure) networks, or \fBibss\fR for  
 2986 \fBIBSS\fR (ad-hoc) networks.  
 2987 .RE

2989 By default, currently all fields but \fBAUTH\fR, \fBBSSID\fR, \fBSSSTYPE\fR are  
 2990 displayed.  
 2991 .RE

2993 .sp  
 2994 .ne 2  
 2995 .na  
 2996 \fB\fB-p\fR, \fB--parseable\fR\fR  
 2997 .ad  
 2998 .sp .6  
 2999 .RS 4n  
 3000 Displays using a stable machine-parseable format. The \fB-o\fR option is  
 3001 required with \fB-p\fR. See "Parseable Output Format", below.  
 3002 .RE

3004 .RE

3006 .sp  
 3007 .ne 2  
 3008 .na  
 3009 \fB\fBdladm show-ether\fR [\fB-x\fR] [\fB-p\fR] \fB-o\fR \fIfield\fR,...  
 3010 [\fIether-link\fR]\fR  
 3011 .ad  
 3012 .sp .6  
 3013 .RS 4n  
 3014 Shows state information either for all physical Ethernet links or for a  
 3015 specified physical Ethernet link.  
 3016 .sp  
 3017 The \fBshow-ether\fR subcommand accepts the following options:  
 3018 .sp  
 3019 .ne 2  
 3020 .na  
 3021 \fB\fB-o\fR \fIfield\fR,..., \fB--output\fR=\fIfield\fR\fR  
 3022 .ad  
 3023 .sp .6  
 3024 .RS 4n  
 3025 A case-insensitive, comma-separated list of output fields to display. The field  
 3026 name must be one of the fields listed below, or the special value \fBall\fR to  
 3027 display all fields. For each link, the following fields can be displayed:  
 3028 .sp  
 3029 .ne 2  
 3030 .na

```

3031 \fB\fBLINK\fR\fR
3032 .ad
3033 .sp .6
3034 .RS 4n
3035 The name of the link being displayed.
3036 .RE

3038 .sp
3039 .ne 2
3040 .na
3041 \fB\fBPTYPE\fR\fR
3042 .ad
3043 .sp .6
3044 .RS 4n
3045 Parameter type, where \fBcurrent\fR indicates the negotiated state of the link,
3046 \fBcapable\fR indicates capabilities supported by the device, \fBadv\fR
3047 indicates the advertised capabilities, and \fBpeeradv\fR indicates the
3048 capabilities advertised by the link-partner.
3049 .RE

3051 .sp
3052 .ne 2
3053 .na
3054 \fB\fBSTATE\fR\fR
3055 .ad
3056 .sp .6
3057 .RS 4n
3058 The state of the link.
3059 .RE

3061 .sp
3062 .ne 2
3063 .na
3064 \fB\fBAUTO\fR\fR
3065 .ad
3066 .sp .6
3067 .RS 4n
3068 A \fByes\fR/\fBno\fR value indicating whether auto-negotiation is advertised.
3069 .RE

3071 .sp
3072 .ne 2
3073 .na
3074 \fB\fBSPEED-DUPLEX\fR\fR
3075 .ad
3076 .sp .6
3077 .RS 4n
3078 Combinations of speed and duplex values available. The units of speed are
3079 encoded with a trailing suffix of \fBG\fR (Gigabits/s) or \fBM\fR (Mb/s).
3080 Duplex values are encoded as \fBF\fR (full-duplex) or \fBH\fR (half-duplex).
3081 .RE

3083 .sp
3084 .ne 2
3085 .na
3086 \fB\fBPAUSE\fR\fR
3087 .ad
3088 .sp .6
3089 .RS 4n
3090 Flow control information. Can be \fBno\fR, indicating no flow control is
3091 available; \fBtx\fR, indicating that the end-point can transmit pause frames,
3092 but ignores any received pause frames; \fBrx\fR, indicating that the end-point
3093 receives and acts upon received pause frames; or \fBbi\fR, indicating
3094 bi-directional flow-control.
3095 .RE

```

```

3097 .sp
3098 .ne 2
3099 .na
3100 \fB\fBREM_FAULT\fR\fR
3101 .ad
3102 .sp .6
3103 .RS 4n
3104 Fault detection information. Valid values are \fBnone\fR or \fBfault\fR.
3105 .RE

3107 By default, all fields except \fBREM_FAULT\fR are displayed for the "current"
3108 \fBPTYPE\fR.
3109 .RE

3111 .sp
3112 .ne 2
3113 .na
3114 \fB\fB-p\fR, \fB--parseable\fR\fR
3115 .ad
3116 .sp .6
3117 .RS 4n
3118 Displays using a stable machine-parseable format. The \fB-o\fR option is
3119 required with \fB-p\fR. See "Parseable Output Format", below.
3120 .RE

3122 .sp
3123 .ne 2
3124 .na
3125 \fB\fB-x\fR, \fB--extended\fR\fR
3126 .ad
3127 .sp .6
3128 .RS 4n
3129 Extended output is displayed for \fBPTYPE\fR values of \fBcurrent\fR,
3130 \fBcapable\fR, \fBadv\fR and \fBpeeradv\fR.
3131 .RE

3133 .RE

3135 .sp
3136 .ne 2
3137 .na
3138 \fB\fBdladm set-linkprop\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-p\fR
3139 \fIprop\fR=\fIvalue\fR[...] \fIlink\fR\fR
3140 .ad
3141 .sp .6
3142 .RS 4n
3143 Sets the values of one or more properties on the link specified. The list of
3144 properties and their possible values depend on the link type, the network
3145 device driver, and networking hardware. These properties can be retrieved using
3146 \fBshow-linkprop\fR.
3147 .sp
3148 .ne 2
3149 .na
3150 \fB\fB-t\fR, \fB--temporary\fR\fR
3151 .ad
3152 .sp .6
3153 .RS 4n
3154 Specifies that the changes are temporary. Temporary changes last until the next
3155 reboot.
3156 .RE

3158 .sp
3159 .ne 2
3160 .na
3161 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3162 .ad

```

```

3163 .sp .6
3164 .RS 4n
3165 See "Options," above.
3166 .RE

3168 .sp
3169 .ne 2
3170 .na
3171 \fB\fB-p\fR \fIprop\fR=\fIvalue\fR[...], \fB--prop\fR
3172 \fIprop\fR=\fIvalue\fR[...]\fR
3173 .ad
3174 .br
3175 .na
3176 \fB\fR
3177 .ad
3178 .sp .6
3179 .RS 4n
3180 A comma-separated list of properties to set to the specified values.
3181 .RE

3183 Note that when the persistent value is set, the temporary value changes to the
3184 same value.
3185 .RE

3187 .sp
3188 .ne 2
3189 .na
3190 \fB\fBdladm reset-linkprop\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-p\fR
3191 \fIprop\fR,...] \fIlink\fR\fR
3192 .ad
3193 .sp .6
3194 .RS 4n
3195 Resets one or more properties to their values on the link specified. Properties
3196 are reset to the values they had at startup. If no properties are specified,
3197 all properties are reset. See \fBshow-linkprop\fR for a description of
3198 properties.
3199 .sp
3200 .ne 2
3201 .na
3202 \fB\fB-t\fR, \fB--temporary\fR\fR
3203 .ad
3204 .sp .6
3205 .RS 4n
3206 Specifies that the resets are temporary. Values are reset to default values.
3207 Temporary resets last until the next reboot.
3208 .RE

3210 .sp
3211 .ne 2
3212 .na
3213 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3214 .ad
3215 .sp .6
3216 .RS 4n
3217 See "Options," above.
3218 .RE

3220 .sp
3221 .ne 2
3222 .na
3223 \fB\fB-p\fR \fIprop, ... \fR, \fB--prop\fR=\fIprop, ... \fR\fR
3224 .ad
3225 .sp .6
3226 .RS 4n
3227 A comma-separated list of properties to reset.
3228 .RE

```

```

3230 Note that when the persistent value is reset, the temporary value changes to
3231 the same value.
3232 .RE

3234 .sp
3235 .ne 2
3236 .na
3237 \fB\fBdladm show-linkprop\fR [\fB-P\fR] [\fB-c\fR] \fB-o\fR
3238 \fIfield\fR[...]\fB-p\fR \fIprop\fR[...]] [\fIlink\fR]\fR
3239 .ad
3240 .sp .6
3241 .RS 4n
3242 Show the current or persistent values of one or more properties, either for all
3243 datalinks or for the specified link. By default, current values are shown. If
3244 no properties are specified, all available link properties are displayed. For
3245 each property, the following fields are displayed:
3246 .sp
3247 .ne 2
3248 .na
3249 \fB\fB-o\fR \fIfield\fR[...], \fB--output\fR=\fIfield\fR\fR
3250 .ad
3251 .sp .6
3252 .RS 4n
3253 A case-insensitive, comma-separated list of output fields to display. The field
3254 name must be one of the fields listed below, or the special value \fBAll\fR to
3255 display all fields. For each link, the following fields can be displayed:
3256 .sp
3257 .ne 2
3258 .na
3259 \fB\fBLINK\fR\fR
3260 .ad
3261 .sp .6
3262 .RS 4n
3263 The name of the datalink.
3264 .RE

3266 .sp
3267 .ne 2
3268 .na
3269 \fB\fBPROPERTY\fR\fR
3270 .ad
3271 .sp .6
3272 .RS 4n
3273 The name of the property.
3274 .RE

3276 .sp
3277 .ne 2
3278 .na
3279 \fB\fBPERM\fR\fR
3280 .ad
3281 .sp .6
3282 .RS 4n
3283 The read/write permissions of the property. The value shown is one of \fBro\fR
3284 or \fBrw\fR.
3285 .RE

3287 .sp
3288 .ne 2
3289 .na
3290 \fB\fBVVALUE\fR\fR
3291 .ad
3292 .sp .6
3293 .RS 4n
3294 The current (or persistent) property value. If the value is not set, it is

```

3295 shown as \fB--\fR. If it is unknown, the value is shown as \fB?\fR. Persistent  
 3296 values that are not set or have been reset will be shown as \fB--\fR and will  
 3297 use the system \fBDEFAULT\fR value (if any).  
 3298 .RE

3300 .sp  
 3301 .ne 2  
 3302 .na  
 3303 \fB\fBDEFAULT\fR\fR  
 3304 .ad  
 3305 .sp .6  
 3306 .RS 4n  
 3307 The default value of the property. If the property has no default value,  
 3308 \fB--\fR is shown.  
 3309 .RE

3311 .sp  
 3312 .ne 2  
 3313 .na  
 3314 \fB\fBPOSSIBLE\fR\fR  
 3315 .ad  
 3316 .sp .6  
 3317 .RS 4n  
 3318 A comma-separated list of the values the property can have. If the values span  
 3319 a numeric range, \fImin\fR - \fImax\fR might be shown as shorthand. If the  
 3320 possible values are unknown or unbounded, \fB--\fR is shown.  
 3321 .RE

3323 The list of properties depends on the link type and network device driver, and  
 3324 the available values for a given property further depends on the underlying  
 3325 network hardware and its state. General link properties are documented in the  
 3326 \fBLINK PROPERTIES\fR section. However, link properties that begin with  
 3327 "\fB\_\fR" (underbar) are specific to a given link or its underlying network  
 3328 device and subject to change or removal. See the appropriate network device  
 3329 driver man page for details.  
 3330 .RE

3332 .sp  
 3333 .ne 2  
 3334 .na  
 3335 \fB\fB-c\fR, \fB--parseable\fR\fR  
 3336 .ad  
 3337 .sp .6  
 3338 .RS 4n  
 3339 Display using a stable machine-parseable format. The \fB-o\fR option is  
 3340 required with this option. See "Parseable Output Format", below.  
 3341 .RE

3343 .sp  
 3344 .ne 2  
 3345 .na  
 3346 \fB\fB-P\fR, \fB--persistent\fR\fR  
 3347 .ad  
 3348 .sp .6  
 3349 .RS 4n  
 3350 Display persistent link property information  
 3351 .RE

3353 .sp  
 3354 .ne 2  
 3355 .na  
 3356 \fB\fB-p\fR \fIprop, ... \fR, \fB--prop\fR=\fIprop, ... \fR\fR  
 3357 .ad  
 3358 .sp .6  
 3359 .RS 4n  
 3360 A comma-separated list of properties to show. See the sections on link

3361 properties following subcommand descriptions.  
 3362 .RE

3364 .RE  
 3366 .sp  
 3367 .ne 2  
 3368 .na  
 3369 \fB\fBdladm create-secobj\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-f\fR  
 3370 \fIfile\fR] \fB-c\fR \fIclass\fR \fIsecobj\fR\fR  
 3371 .ad  
 3372 .sp .6  
 3373 .RS 4n  
 3374 Create a secure object named \fIsecobj\fR in the specified \fIclass\fR to be  
 3375 later used as a WEP or WPA key in connecting to an encrypted network. The value  
 3376 of the secure object can either be provided interactively or read from a file.  
 3377 The sequence of interactive prompts and the file format depends on the class of  
 3378 the secure object.  
 3379 .sp  
 3380 Currently, the classes \fBWep\fR and \fBwpa\fR are supported. The \fBWEP\fR  
 3381 (Wired Equivalent Privacy) key can be either 5 or 13 bytes long. It can be  
 3382 provided either as an \fBASCII\fR or hexadecimal string -- thus, \fB12345\fR  
 3383 and \fB0x3132333435\fR are equivalent 5-byte keys (the \fB0\fR prefix can be  
 3384 omitted). A file containing a \fBWEP\fR key must consist of a single line using  
 3385 either \fBWEP\fR key format. The WPA (Wi-Fi Protected Access) key must be  
 3386 provided as an ASCII string with a length between 8 and 63 bytes.  
 3387 .sp  
 3388 This subcommand is only usable by users or roles that belong to the "Network  
 3389 Link Security" \fBRBAC\fR profile.  
 3390 .sp  
 3391 .ne 2  
 3392 .na  
 3393 \fB\fB-c\fR \fIclass\fR, \fB--class\fR=\fIclass\fR\fR  
 3394 .ad  
 3395 .sp .6  
 3396 .RS 4n  
 3397 \fIclass\fR can be \fBWep\fR or \fBwpa\fR. See preceding discussion.  
 3398 .RE

3400 .sp  
 3401 .ne 2  
 3402 .na  
 3403 \fB\fB-t\fR, \fB--temporary\fR\fR  
 3404 .ad  
 3405 .sp .6  
 3406 .RS 4n  
 3407 Specifies that the creation is temporary. Temporary creation last until the  
 3408 next reboot.  
 3409 .RE

3411 .sp  
 3412 .ne 2  
 3413 .na  
 3414 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR  
 3415 .ad  
 3416 .sp .6  
 3417 .RS 4n  
 3418 See "Options," above.  
 3419 .RE

3421 .sp  
 3422 .ne 2  
 3423 .na  
 3424 \fB\fB-f\fR \fIfile\fR, \fB--file\fR=\fIfile\fR\fR  
 3425 .ad  
 3426 .sp .6

```

3427 .RS 4n
3428 Specifies a file that should be used to obtain the secure object's value. The
3429 format of this file depends on the secure object class. See the \fBEXAMPLES\fR
3430 section for an example of using this option to set a \fBWEP\fR key.
3431 .RE

3433 .RE

3435 .sp
3436 .ne 2
3437 .na
3438 \fB\fBdladm delete-secobj\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR]
3439 [\fIsecobj\fR[...]\fR
3440 .ad
3441 .sp .6
3442 .RS 4n
3443 Delete one or more specified secure objects. This subcommand is only usable by
3444 users or roles that belong to the "Network Link Security" \fBRBAC\fR profile.
3445 .sp
3446 .ne 2
3447 .na
3448 \fB\fB-t\fR, \fB--temporary\fR\fR
3449 .ad
3450 .sp .6
3451 .RS 4n
3452 Specifies that the deletions are temporary. Temporary deletions last until the
3453 next reboot.
3454 .RE

3456 .sp
3457 .ne 2
3458 .na
3459 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3460 .ad
3461 .sp .6
3462 .RS 4n
3463 See "Options," above.
3464 .RE

3466 .RE

3468 .sp
3469 .ne 2
3470 .na
3471 \fB\fBdladm show-secobj\fR [\fB-P\fR] [\fB-o\fR \fIfield\fR[...]]
3472 [\fIsecobj\fR,...]\fR
3473 .ad
3474 .sp .6
3475 .RS 4n
3476 Show current or persistent secure object information. If one or more secure
3477 objects are specified, then information for each is displayed. Otherwise, all
3478 current or persistent secure objects are displayed.
3479 .sp
3480 By default, current secure objects are displayed, which are all secure objects
3481 that have either been persistently created and not temporarily deleted, or
3482 temporarily created.
3483 .sp
3484 For security reasons, it is not possible to show the value of a secure object.
3485 .sp
3486 .ne 2
3487 .na
3488 \fB\fB-o\fR \fIfield\fR[...], \fB--output\fR=\fIfield\fR[...]\fR
3489 .ad
3490 .sp .6
3491 .RS 4n
3492 A case-insensitive, comma-separated list of output fields to display. The field

```

```

3493 name must be one of the fields listed below. For displayed secure object, the
3494 following fields can be shown:
3495 .sp
3496 .ne 2
3497 .na
3498 \fB\fBOBJECT\fR\fR
3499 .ad
3500 .sp .6
3501 .RS 4n
3502 The name of the secure object.
3503 .RE

3505 .sp
3506 .ne 2
3507 .na
3508 \fB\fBCCLASS\fR\fR
3509 .ad
3510 .sp .6
3511 .RS 4n
3512 The class of the secure object.
3513 .RE

3515 .RE

3517 .sp
3518 .ne 2
3519 .na
3520 \fB\fB-p\fR, \fB--parseable\fR\fR
3521 .ad
3522 .sp .6
3523 .RS 4n
3524 Display using a stable machine-parseable format. The \fB-o\fR option is
3525 required with \fB-p\fR. See "Parseable Output Format", below.
3526 .RE

3528 .sp
3529 .ne 2
3530 .na
3531 \fB\fB-P\fR, \fB--persistent\fR\fR
3532 .ad
3533 .sp .6
3534 .RS 4n
3535 Display persistent secure object information
3536 .RE

3538 .RE

3540 .sp
3541 .ne 2
3542 .na
3543 \fB\fBdladm create-vnic\fR [\fB-t\fR] \fB-1\fR \fIlink\fR [\fB-R\fR
3544 \fIroot-dir\fR] [\fB-m\fR \fIvalue\fR | auto | {factory [\fB-n\fR
3545 \fIslot-identifier\fR]} | {random [\fB-r\fR \fIprefix\fR]}] [\fB-v\fR
3546 \fIvlan-id\fR] [\fB-p\fR \fIprop\fR=\fIvalue\fR[...]] \fIvnic-link\fR\fR
3547 .ad
3548 .sp .6
3549 .RS 4n
3550 Create a VNIC with name \fIvnic-link\fR over the specified link.
3551 .sp
3552 .ne 2
3553 .na
3554 \fB\fB-t\fR, \fB--temporary\fR\fR
3555 .ad
3556 .sp .6
3557 .RS 4n
3558 Specifies that the VNIC is temporary. Temporary VNICs last until the next

```

```

3559 reboot.
3560 .RE

3562 .sp
3563 .ne 2
3564 .na
3565 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3566 .ad
3567 .sp .6
3568 .RS 4n
3569 See "Options," above.
3570 .RE

3572 .sp
3573 .ne 2
3574 .na
3575 \fB\fB-1\fR \fIlink\fR, \fB--link\fR=\fIlink\fR\fR
3576 .ad
3577 .sp .6
3578 .RS 4n
3579 \fIlink\fR can be a physical link or an \fBetherstub\fR.
3580 .RE

3582 .sp
3583 .ne 2
3584 .na
3585 \fB\fB-m\fR \fIvalue\fR | \fIkeyword\fR, \fB--mac-address\fR=\fIvalue\fR |
3586 \fIkeyword\fR\fR
3587 .ad
3588 .sp .6
3589 .RS 4n
3590 Sets the VNIC's MAC address based on the specified value or keyword. If
3591 \fIvalue\fR is not a keyword, it is interpreted as a unicast MAC address, which
3592 must be valid for the underlying NIC. The following special keywords can be
3593 used:
3594 .sp
3595 .ne 2
3596 .na
3597 \fBfactory [\fB-n\fR \fIslot-identifier\fR],\fR
3598 .ad
3599 .br
3600 .na
3601 \fBfactory [\fB--slot\fR=\fIslot-identifier\fR]\fR
3602 .ad
3603 .sp .6
3604 .RS 4n
3605 Assign a factory MAC address to the VNIC. When a factory MAC address is
3606 requested, \fB-m\fR can be combined with the \fB-n\fR option to specify a MAC
3607 address slot to be used. If \fB-n\fR is not specified, the system will choose
3608 the next available factory MAC address. The \fB-m\fR option of the
3609 \fBshow-phys\fR subcommand can be used to display the list of factory MAC
3610 addresses, their slot identifiers, and their availability.
3611 .RE

3613 .sp
3614 .ne 2
3615 .na
3616 \fB\fR
3617 .ad
3618 .br
3619 .na
3620 \fBrandom [\fB-r\fR \fIprefix\fR],\fR
3621 .ad
3622 .br
3623 .na
3624 \fBrandom [\fB--mac-prefix\fR=\fIprefix\fR]\fR

```

```

3625 .ad
3626 .sp .6
3627 .RS 4n
3628 Assign a random MAC address to the VNIC. A default prefix consisting of a valid
3629 IEEE OUI with the local bit set will be used. That prefix can be overridden
3630 with the \fB-r\fR option.
3631 .RE

3633 .sp
3634 .ne 2
3635 .na
3636 \fBauto\fR
3637 .ad
3638 .sp .6
3639 .RS 4n
3640 Try and use a factory MAC address first. If none is available, assign a random
3641 MAC address. \fBauto\fR is the default action if the \fB-m\fR option is not
3642 specified.
3643 .RE

3645 .sp
3646 .ne 2
3647 .na
3648 \fB\fB-v\fR \fIvlan-id\fR\fR
3649 .ad
3650 .sp .6
3651 .RS 4n
3652 Enable VLAN tagging for this VNIC. The VLAN tag will have id \fIvlan-id\fR.
3653 .RE

3655 .RE

3657 .sp
3658 .ne 2
3659 .na
3660 \fB\fB-p\fR \fIprop\fR=\fIvalue\fR,..., \fB--prop\fR
3661 \fIprop\fR=\fIvalue\fR,...\fR
3662 .ad
3663 .sp .6
3664 .RS 4n
3665 A comma-separated list of properties to set to the specified values.
3666 .RE

3668 .RE

3670 .sp
3671 .ne 2
3672 .na
3673 \fB\fBdadm delete-vnic\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR]
3674 \fBnic-link\fR\fR
3675 .ad
3676 .sp .6
3677 .RS 4n
3678 Deletes the specified VNIC.
3679 .sp
3680 .ne 2
3681 .na
3682 \fB\fB-t\fR, \fB--temporary\fR\fR
3683 .ad
3684 .sp .6
3685 .RS 4n
3686 Specifies that the deletion is temporary. Temporary deletions last until the
3687 next reboot.
3688 .RE

3690 .sp

```

```

3691 .ne 2
3692 .na
3693 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3694 .ad
3695 .sp .6
3696 .RS 4n
3697 See "Options," above.
3698 .RE

3700 .RE

3702 .sp
3703 .ne 2
3704 .na
3705 \fB\fBdadm show-vnic\fR [\fB-pp\fR] [\fB-s\fR [\fB-i\fR \fIinterval\fR]]
3706 [\fB-o\fR \fIfield\fR[,...]] [\fB-l\fR \fIlink\fR] [\fIvnic-link\fR]\fR
3707 .ad
3708 .sp .6
3709 .RS 4n
3710 Show VNIC configuration information (the default) or statistics, for all VNICs,
3711 all VNICs on a link, or only the specified \fIvnic-link\fR.
3712 .sp
3713 .ne 2
3714 .na
3715 \fB\fB-o\fR \fIfield\fR[,...], \fB--output\fR=\fIfield\fR[,...]\fR
3716 .ad
3717 .sp .6
3718 .RS 4n
3719 A case-insensitive, comma-separated list of output fields to display. The field
3720 name must be one of the fields listed below. The field name must be one of the
3721 fields listed below, or the special value \fBall\fR to display all fields. By
3722 default (without \fB-o\fR), \fBshow-vnic\fR displays all fields.
3723 .sp
3724 .ne 2
3725 .na
3726 \fB\fBLINK\fR\fR
3727 .ad
3728 .sp .6
3729 .RS 4n
3730 The name of the VNIC.
3731 .RE

3733 .sp
3734 .ne 2
3735 .na
3736 \fB\fBOVER\fR\fR
3737 .ad
3738 .sp .6
3739 .RS 4n
3740 The name of the physical link over which this VNIC is configured.
3741 .RE

3743 .sp
3744 .ne 2
3745 .na
3746 \fB\fBSPEED\fR\fR
3747 .ad
3748 .sp .6
3749 .RS 4n
3750 The maximum speed of the VNIC, in megabits per second.
3751 .RE

3753 .sp
3754 .ne 2
3755 .na
3756 \fB\fBMACADDRESS\fR\fR

```

```

3757 .ad
3758 .sp .6
3759 .RS 4n
3760 MAC address of the VNIC.
3761 .RE

3763 .sp
3764 .ne 2
3765 .na
3766 \fB\fBMACADDRTYPE\fR\fR
3767 .ad
3768 .sp .6
3769 .RS 4n
3770 MAC address type of the VNIC. \fBdadm\fR distinguishes among the following MAC
3771 address types:
3772 .sp
3773 .ne 2
3774 .na
3775 \fB\fBrandom\fR\fR
3776 .ad
3777 .sp .6
3778 .RS 4n
3779 A random address assigned to the VNIC.
3780 .RE

3782 .sp
3783 .ne 2
3784 .na
3785 \fB\fBfactory\fR\fR
3786 .ad
3787 .sp .6
3788 .RS 4n
3789 A factory MAC address used by the VNIC.
3790 .RE

3792 .RE

3794 .RE

3796 .sp
3797 .ne 2
3798 .na
3799 \fB\fB-p\fR, \fB--parseable\fR\fR
3800 .ad
3801 .sp .6
3802 .RS 4n
3803 Display using a stable machine-parseable format. The \fB-o\fR option is
3804 required with \fB-p\fR. See "Parseable Output Format", below.
3805 .RE

3807 .sp
3808 .ne 2
3809 .na
3810 \fB\fB-P\fR, \fB--persistent\fR\fR
3811 .ad
3812 .sp .6
3813 .RS 4n
3814 Display the persistent VNIC configuration.
3815 .RE

3817 .sp
3818 .ne 2
3819 .na
3820 \fB\fB-s\fR, \fB--statistics\fR\fR
3821 .ad
3822 .sp .6

```

```

3823 .RS 4n
3824 Displays VNIC statistics.
3825 .RE

3827 .sp
3828 .ne 2
3829 .na
3830 \fB\fB-i\fR \fIinterval\fR, \fB--interval\fR=\fIinterval\fR\fR
3831 .ad
3832 .sp .6
3833 .RS 4n
3834 Used with the \fB-s\fR option to specify an interval, in seconds, at which
3835 statistics should be displayed. If this option is not specified, statistics
3836 will be displayed only once.
3837 .RE

3839 .sp
3840 .ne 2
3841 .na
3842 \fB\fB-l\fR \fIlink\fR, \fB--link\fR=\fIlink\fR\fR
3843 .ad
3844 .sp .6
3845 .RS 4n
3846 Display information for all VNICs on the named link.
3847 .RE

3849 .RE

3851 .sp
3852 .ne 2
3853 .na
3854 \fB\fR
3855 .ad
3856 .br
3857 .na
3858 \fB\fBdadm create-etherstub\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR]
3859 \fIetherstub\fR\fR
3860 .ad
3861 .sp .6
3862 .RS 4n
3863 Create an etherstub with the specified name.
3864 .sp
3865 .ne 2
3866 .na
3867 \fB\fB-t\fR, \fB--temporary\fR\fR
3868 .ad
3869 .sp .6
3870 .RS 4n
3871 Specifies that the etherstub is temporary. Temporary etherstubs do not persist
3872 across reboots.
3873 .RE

3875 .sp
3876 .ne 2
3877 .na
3878 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3879 .ad
3880 .sp .6
3881 .RS 4n
3882 See "Options," above.
3883 .RE

3885 VNICs can be created on top of etherstubs instead of physical NICs. As with
3886 physical NICs, such a creation causes the stack to implicitly create a virtual
3887 switch between the VNICs created on top of the same etherstub.
3888 .RE

```

```

3890 .sp
3891 .ne 2
3892 .na
3893 \fB\fR
3894 .ad
3895 .br
3896 .na
3897 \fB\fBdadm delete-etherstub\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR]
3898 \fIetherstub\fR\fR
3899 .ad
3900 .sp .6
3901 .RS 4n
3902 Delete the specified etherstub.
3903 .sp
3904 .ne 2
3905 .na
3906 \fB\fB-t\fR, \fB--temporary\fR\fR
3907 .ad
3908 .sp .6
3909 .RS 4n
3910 Specifies that the deletion is temporary. Temporary deletions last until the
3911 next reboot.
3912 .RE

3914 .sp
3915 .ne 2
3916 .na
3917 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3918 .ad
3919 .sp .6
3920 .RS 4n
3921 See "Options," above.
3922 .RE

3924 .RE

3926 .sp
3927 .ne 2
3928 .na
3929 \fB\fBdadm show-etherstub\fR [\fIetherstub\fR]\fR
3930 .ad
3931 .sp .6
3932 .RS 4n
3933 Show all configured etherstubs by default, or the specified etherstub if
3934 \fIetherstub\fR is specified.
3935 .RE

3937 .sp
3938 .ne 2
3939 .na
3940 \fB\fBdadm create-iptun\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] \fB-T\fR
3941 \fItype\fR [\fB-s\fR \fItsrc\fR] [\fB-d\fR \fItdst\fR] \fIiptun-link\fR\fR
3942 .ad
3943 .sp .6
3944 .RS 4n
3945 Create an IP tunnel link named \fIiptun-link\fR. Such links can additionally be
3946 protected with IPsec using \fBipsecconf\fR(1M).
3947 .sp
3948 An IP tunnel is conceptually comprised of two parts: a virtual link between two
3949 or more IP nodes, and an IP interface above this link that allows the system to
3950 transmit and receive IP packets encapsulated by the underlying link. This
3951 subcommand creates a virtual link. The \fBifconfig\fR(1M) command is used to
3952 configure IP interfaces above the link.
3953 .sp
3954 .ne 2

```

```

3955 .na
3956 \fB\fB-t\fR, \fB--temporary\fR\fR
3957 .ad
3958 .sp .6
3959 .RS 4n
3960 Specifies that the IP tunnel link is temporary. Temporary tunnels last until
3961 the next reboot.
3962 .RE

3964 .sp
3965 .ne 2
3966 .na
3967 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
3968 .ad
3969 .sp .6
3970 .RS 4n
3971 See "Options," above.
3972 .RE

3974 .sp
3975 .ne 2
3976 .na
3977 \fB\fB-T\fR \fItype\fR, \fB--tunnel-type\fR=\fItype\fR\fR
3978 .ad
3979 .sp .6
3980 .RS 4n
3981 Specifies the type of tunnel to be created. The type must be one of the
3982 following:
3983 .sp
3984 .ne 2
3985 .na
3986 \fB\fBipv4\fR\fR
3987 .ad
3988 .sp .6
3989 .RS 4n
3990 A point-to-point, IP-over-IP tunnel between two IPv4 nodes. This type of tunnel
3991 requires IPv4 source and destination addresses to function. IPv4 and IPv6
3992 interfaces can be plumbed above such a tunnel to create IPv4-over-IPv4 and
3993 IPv6-over-IPv4 tunneling configurations.
3994 .RE

3996 .sp
3997 .ne 2
3998 .na
3999 \fB\fBipv6\fR\fR
4000 .ad
4001 .sp .6
4002 .RS 4n
4003 A point-to-point, IP-over-IP tunnel between two IPv6 nodes as defined in IETF
4004 RFC 2473. This type of tunnel requires IPv6 source and destination addresses to
4005 function. IPv4 and IPv6 interfaces can be plumbed above such a tunnel to create
4006 IPv4-over-IPv6 and IPv6-over-IPv6 tunneling configurations.
4007 .RE

4009 .sp
4010 .ne 2
4011 .na
4012 \fB\fB6to4\fR\fR
4013 .ad
4014 .sp .6
4015 .RS 4n
4016 A 6to4, point-to-multipoint tunnel as defined in IETF RFC 3056. This type of
4017 tunnel requires an IPv4 source address to function. An IPv6 interface is
4018 plumbed on such a tunnel link to configure a 6to4 router.
4019 .RE

```

```

4021 .RE
4023 .sp
4024 .ne 2
4025 .na
4026 \fB\fB-s\fR \fIitsrc\fR, \fB--tunnel-src\fR=\fIitsrc\fR\fR
4027 .ad
4028 .sp .6
4029 .RS 4n
4030 Literal IP address or hostname corresponding to the tunnel source. If a
4031 hostname is specified, it will be resolved to IP addresses, and one of those IP
4032 addresses will be used as the tunnel source. Because IP tunnels are created
4033 before naming services have been brought online during the boot process, it is
4034 important that any hostname used be included in \fB/etc/hosts\fR.
4035 .RE

4037 .sp
4038 .ne 2
4039 .na
4040 \fB\fB-d\fR \fItdst\fR, \fB--tunnel-dst\fR=\fItdst\fR\fR
4041 .ad
4042 .sp .6
4043 .RS 4n
4044 Literal IP address or hostname corresponding to the tunnel destination.
4045 .RE

4047 .RE
4049 .sp
4050 .ne 2
4051 .na
4052 \fB\fBdladm modify-iptun\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR] [\fB-s\fR
4053 \fIitsrc\fR] [\fB-d\fR \fItdst\fR] \fIiptun-link\fR\fR
4054 .ad
4055 .sp .6
4056 .RS 4n
4057 Modify the parameters of the specified IP tunnel.
4058 .sp
4059 .ne 2
4060 .na
4061 \fB\fB-t\fR, \fB--temporary\fR\fR
4062 .ad
4063 .sp .6
4064 .RS 4n
4065 Specifies that the modification is temporary. Temporary modifications last
4066 until the next reboot.
4067 .RE

4069 .sp
4070 .ne 2
4071 .na
4072 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
4073 .ad
4074 .sp .6
4075 .RS 4n
4076 See "Options," above.
4077 .RE

4079 .sp
4080 .ne 2
4081 .na
4082 \fB\fB-s\fR \fIitsrc\fR, \fB--tunnel-src\fR=\fIitsrc\fR\fR
4083 .ad
4084 .sp .6
4085 .RS 4n
4086 Specifies a new tunnel source address. See \fBcreate-iptun\fR for a

```

```

4087 description.
4088 .RE

4090 .sp
4091 .ne 2
4092 .na
4093 \fB\fB-d\fR \fItdst\fR, \fB--tunnel-dst\fR=\fItdst\fR\fR
4094 .ad
4095 .sp .6
4096 .RS 4n
4097 Specifies a new tunnel destination address. See \fBcreate-iptun\fR for a
4098 description.
4099 .RE

4101 .RE

4103 .sp
4104 .ne 2
4105 .na
4106 \fB\fBdadm delete-iptun\fR [\fB-t\fR] [\fB-R\fR \fIroot-dir\fR]
4107 \fIiptun-link\fR\fR
4108 .ad
4109 .sp .6
4110 .RS 4n
4111 Delete the specified IP tunnel link.
4112 .sp
4113 .ne 2
4114 .na
4115 \fB\fB-t\fR, \fB--temporary\fR\fR
4116 .ad
4117 .sp .6
4118 .RS 4n
4119 Specifies that the deletion is temporary. Temporary deletions last until the
4120 next reboot.
4121 .RE

4123 .sp
4124 .ne 2
4125 .na
4126 \fB\fB-R\fR \fIroot-dir\fR, \fB--root-dir\fR=\fIroot-dir\fR\fR
4127 .ad
4128 .sp .6
4129 .RS 4n
4130 See "Options," above.
4131 .RE

4133 .RE

4135 .sp
4136 .ne 2
4137 .na
4138 \fB\fBdadm show-iptun\fR [\fB-P\fR] [[\fB-p\fR] \fB-o\fR \fIfield\fR[...]]
4139 [\fIiptun-link\fR]\fR
4140 .ad
4141 .sp .6
4142 .RS 4n
4143 Show IP tunnel link configuration for a single IP tunnel or all IP tunnels.
4144 .sp
4145 .ne 2
4146 .na
4147 \fB\fB-P\fR, \fB--persistent\fR\fR
4148 .ad
4149 .sp .6
4150 .RS 4n
4151 Display the persistent IP tunnel configuration.
4152 .RE

```

```

4154 .sp
4155 .ne 2
4156 .na
4157 \fB\fB-p\fR, \fB--parseable\fR\fR
4158 .ad
4159 .sp .6
4160 .RS 4n
4161 Display using a stable machine-parseable format. The -o option is required with
4162 -p. See "Parseable Output Format", below.
4163 .RE

4165 .sp
4166 .ne 2
4167 .na
4168 \fB\fB-o\fR \fIfield\fR[...], \fB--output\fR=\fIfield\fR[...]\fR
4169 .ad
4170 .sp .6
4171 .RS 4n
4172 A case-insensitive, comma-separated list of output fields to display. The field
4173 name must be one of the fields listed below, or the special value \fBall\fR, to
4174 display all fields. By default (without \fB-o\fR), \fBshow-iptun\fR displays
4175 all fields.
4176 .sp
4177 .ne 2
4178 .na
4179 \fB\fBLINK\fR\fR
4180 .ad
4181 .sp .6
4182 .RS 4n
4183 The name of the IP tunnel link.
4184 .RE

4186 .sp
4187 .ne 2
4188 .na
4189 \fB\fBTYPE\fR\fR
4190 .ad
4191 .sp .6
4192 .RS 4n
4193 Type of tunnel as specified by the \fB-T\fR option of \fBcreate-iptun\fR.
4194 .RE

4196 .sp
4197 .ne 2
4198 .na
4199 \fB\fBFLAGS\fR\fR
4200 .ad
4201 .sp .6
4202 .RS 4n
4203 A set of flags associated with the IP tunnel link. Possible flags are:
4204 .sp
4205 .ne 2
4206 .na
4207 \fB\fBs\fR\fR
4208 .ad
4209 .sp .6
4210 .RS 4n
4211 The IP tunnel link is protected by IPsec policy. To display the IPsec policy
4212 associated with the tunnel link, enter:
4213 .sp
4214 .in +2
4215 .nf
4216 # \fBipsecconf -ln -i \fItunnel-link\fR\fR
4217 .fi
4218 .in -2

```

```

4219 .sp
4221 See \fBipsecconf\fR(1M) for more details on how to configure IPsec policy.
4222 .RE

4224 .sp
4225 .ne 2
4226 .na
4227 \fB\fBi\fR
4228 .ad
4229 .sp .6
4230 .RS 4n
4231 The IP tunnel link was implicitly created with \fBifconfig\fR(1M), and will be
4232 automatically deleted when it is no longer referenced (that is, when the last
4233 IP interface over the tunnel is unplumbed). See \fBifconfig\fR(1M) for details
4234 on implicit tunnel creation.
4235 .RE

4237 .RE

4239 .sp
4240 .ne 2
4241 .na
4242 \fB\fBSOURCE\fR\fR
4243 .ad
4244 .sp .6
4245 .RS 4n
4246 The tunnel source address.
4247 .RE

4249 .sp
4250 .ne 2
4251 .na
4252 \fB\fBDESTINATION\fR\fR
4253 .ad
4254 .sp .6
4255 .RS 4n
4256 The tunnel destination address.
4257 .RE

4259 .RE

4261 .RE

4263 .sp
4264 .ne 2
4265 .na
4266 \fB\fBdladm show-usage\fR [\fB-a\fR] \fB-f\fR \fIfilename\fR [\fB-p\fR
4267 \fIplotfile\fR \fB-F\fR \fIformat\fR] [\fB-s\fR \fItime\fR] [\fB-e\fR
4268 \fItime\fR] [\filink\fR]\fR
4269 .ad
4270 .sp .6
4271 .RS 4n
4272 Show the historical network usage from a stored extended accounting file.
4273 Configuration and enabling of network accounting through \fBacctadm\fR(1M) is
4274 required. The default output will be the summary of network usage for the
4275 entire period of time in which extended accounting was enabled.
4276 .sp
4277 .ne 2
4278 .na
4279 \fB\fB-a\fR\fR
4280 .ad
4281 .sp .6
4282 .RS 4n
4283 Display all historical network usage for the specified period of time during
4284 which extended accounting is enabled. This includes the usage information for

```

```

4285 the links that have already been deleted.
4286 .RE

4288 .sp
4289 .ne 2
4290 .na
4291 \fB\fB-f\fR \fIfilename\fR, \fB--file\fR=\fIfilename\fR\fR
4292 .ad
4293 .sp .6
4294 .RS 4n
4295 Read extended accounting records of network usage from \fIfilename\fR.
4296 .RE

4298 .sp
4299 .ne 2
4300 .na
4301 \fB\fB-F\fR \fIformat\fR, \fB--format\fR=\fIformat\fR\fR
4302 .ad
4303 .sp .6
4304 .RS 4n
4305 Specifies the format of \fIplotfile\fR that is specified by the \fB-p\fR
4306 option. As of this release, \fBgnuplot\fR is the only supported format.
4307 .RE

4309 .sp
4310 .ne 2
4311 .na
4312 \fB\fB-p\fR \fIplotfile\fR, \fB--plot\fR=\fIplotfile\fR\fR
4313 .ad
4314 .sp .6
4315 .RS 4n
4316 Write network usage data to a file of the format specified by the \fB-F\fR
4317 option, which is required.
4318 .RE

4320 .sp
4321 .ne 2
4322 .na
4323 \fB\fB-s\fR \fItime\fR, \fB--start\fR=\fItime\fR\fR
4324 .ad
4325 .br
4326 .na
4327 \fB\fB-e\fR \fItime\fR, \fB--stop\fR=\fItime\fR\fR
4328 .ad
4329 .sp .6
4330 .RS 4n
4331 Start and stop times for data display. Time is in the format
4332 \fIMMM\fR/\fIDDD\fR/\fYYYY\fR,\fIhh\fR:\fImm\fR:\fIss\fR.
4333 .RE

4335 .sp
4336 .ne 2
4337 .na
4338 \fB\fIlink\fR\fR
4339 .ad
4340 .sp .6
4341 .RS 4n
4342 If specified, display the network usage only for the named link. Otherwise,
4343 display network usage for all links.
4344 .RE

4346 .RE

4348 .SS "Parseable Output Format"
4349 .sp
4350 .LP

```

```

4351 Many \fBdadm\fR subcommands have an option that displays output in a
4352 machine-parseable format. The output format is one or more lines of colon
4353 (\fB:\fR) delimited fields. The fields displayed are specific to the subcommand
4354 used and are listed under the entry for the \fB-o\fR option for a given
4355 subcommand. Output includes only those fields requested by means of the
4356 \fB-o\fR option, in the order requested.
4357 .sp
4358 .LP
4359 When you request multiple fields, any literal colon characters are escaped by a
4360 backslash (\fB\:\fR) before being output. Similarly, literal backslash
4361 characters will also be escaped (\fB\\:\fR). This escape format is parseable
4362 by using shell \fBread\fR(1) functions with the environment variable
4363 \fBFIFS=\fR (see \fBEXAMPLES\fR, below). Note that escaping is not done when
4364 you request only a single field.
4365 .SS "General Link Properties"
4366 .sp
4367 .LP
4368 The following general link properties are supported:
4369 .sp
4370 .ne 2
4371 .na
4372 \fB\fBautopush\fR\fR
4373 .ad
4374 .sp .6
4375 .RS 4n
4376 Specifies the set of STREAMS modules to push on the stream associated with a
4377 link when its DLPI device is opened. It is a space-delimited list of modules.
4378 .sp
4379 The optional special character sequence \fB[anchor]\fR indicates that a STREAMS
4380 anchor should be placed on the stream at the module previously specified in the
4381 list. It is an error to specify more than one anchor or to have an anchor first
4382 in the list.
4383 .sp
4384 The \fBautopush\fR property is preferred over the more general
4385 \fBautopush\fR(1M) command.
4386 .RE

4388 .sp
4389 .ne 2
4390 .na
4391 \fB\fBcpus\fR\fR
4392 .ad
4393 .sp .6
4394 .RS 4n
4395 Bind the processing of packets for a given data link to a processor or a set of
4396 processors. The value can be a comma-separated list of one or more processor
4397 ids. If the list consists of more than one processor, the processing will
4398 spread out to all the processors. Connection to processor affinity and packet
4399 ordering for any individual connection will be maintained.
4400 .sp
4401 The processor or set of processors are not exclusively reserved for the link.
4402 Only the kernel threads and interrupts associated with processing of the link
4403 are bound to the processor or the set of processors specified. In case it is
4404 desired that processors be dedicated to the link, \fBpsrset\fR(1M) can be used
4405 to create a processor set and then specifying the processors from the processor
4406 set to bind the link to.
4407 .sp
4408 If the link was already bound to processor or set of processors due to a
4409 previous operation, the binding will be removed and the new set of processors
4410 will be used instead.
4411 .sp
4412 The default is no CPU binding, which is to say that the processing of packets
4413 is not bound to any specific processor or processor set.
4414 .RE

4416 .sp

```

```

4417 .ne 2
4418 .na
4419 \fB\fBlearn_limit\fR\fR
4420 .ad
4421 .sp .6
4422 .RS 4n
4423 Limits the number of new or changed MAC sources to be learned over a bridge
4424 link. When the number exceeds this value, learning on that link is temporarily
4425 disabled. Only non-VLAN, non-VNIC type links have this property.
4426 .sp
4427 The default value is \fB1000\fR. Valid values are greater or equal to 0.
4428 .RE

4430 .sp
4431 .ne 2
4432 .na
4433 \fB\fBlearn_decay\fR\fR
4434 .ad
4435 .sp .6
4436 .RS 4n
4437 Specifies the decay rate for source changes limited by \fBlearn_limit\fR. This
4438 number is subtracted from the counter for a bridge link every 5 seconds. Only
4439 non-VLAN, non-VNIC type links have this property.
4440 .sp
4441 The default value is \fB200\fR. Valid values are greater or equal to 0.
4442 .RE

4444 .sp
4445 .ne 2
4446 .na
4447 \fB\fBmac-address\fR\fR
4448 .ad
4449 .sp .6
4450 .RS 4n
4451 The MAC address of the link. The default value is the factory MAC address.
4452 .RE

4454 .sp
4455 .ne 2
4456 .na
4457 #endif /* ! codereview */
4458 \fB\fBmaxbw\fR\fR
4459 .ad
4460 .sp .6
4461 .RS 4n
4462 Sets the full duplex bandwidth for the link. The bandwidth is specified as an
4463 integer with one of the scale suffixes (\fBK\fR, \fBM\fR, or \fBG\fR for Kbps,
4464 Mbps, and Gbps). If no units are specified, the input value will be read as
4465 Mbps. The default is no bandwidth limit.
4466 .RE

4468 .sp
4469 .ne 2
4470 .na
4471 \fB\fBpriority\fR\fR
4472 .ad
4473 .sp .6
4474 .RS 4n
4475 Sets the relative priority for the link. The value can be given as one of the
4476 tokens \fBhigh\fR, \fBmedium\fR, or \fBlow\fR. The default is \fBhigh\fR.
4477 .RE

4479 .sp
4480 .ne 2
4481 .na
4482 \fB\fBstp\fR\fR

```

```

4483 .ad
4484 .sp .6
4485 .RS 4n
4486 Enables or disables Spanning Tree Protocol on a bridge link. Setting this value
4487 to \fB\fR disables Spanning Tree, and puts the link into forwarding mode with
4488 BPDU guarding enabled. This mode is appropriate for point-to-point links
4489 connected only to end nodes. Only non-VLAN, non-VNIC type links have this
4490 property. The default value is \fB1\fR, to enable STP.
4491 .RE

4493 .sp
4494 .ne 2
4495 .na
4496 \fB\fBforward\fR\fR
4497 .ad
4498 .sp .6
4499 .RS 4n
4500 Enables or disables forwarding for a VLAN. Setting this value to \fB0\fR
4501 disables bridge forwarding for a VLAN link. Disabling bridge forwarding removes
4502 that VLAN from the "allowed set" for the bridge. The default value is \fB1\fR,
4503 to enable bridge forwarding for configured VLANs.
4504 .RE

4506 .sp
4507 .ne 2
4508 .na
4509 \fB\fBdefault_tag\fR\fR
4510 .ad
4511 .sp .6
4512 .RS 4n
4513 Sets the default VLAN ID that is assumed for untagged packets sent to and
4514 received from this link. Only non-VLAN, non-VNIC type links have this property.
4515 Setting this value to \fB0\fR disables the bridge forwarding of untagged
4516 packets to and from the port. The default value is \fBVLAN ID 1\fR. Valid
4517 values values are from 0 to 4094.
4518 .RE

4520 .sp
4521 .ne 2
4522 .na
4523 \fB\fBstp_priority\fR\fR
4524 .ad
4525 .sp .6
4526 .RS 4n
4527 Sets the STP and RSTP Port Priority value, which is used to determine the
4528 preferred root port on a bridge. Lower numerical values are higher priority.
4529 The default value is \fB128\fR. Valid values range from 0 to 255.
4530 .RE

4532 .sp
4533 .ne 2
4534 .na
4535 \fB\fBstp_cost\fR\fR
4536 .ad
4537 .sp .6
4538 .RS 4n
4539 Sets the STP and RSTP cost for using the link. The default value is \fBauto\fR,
4540 which sets the cost based on link speed, using \fB100\fR for 10Mbps, \fB19\fR
4541 for 100Mbps, \fB4\fR for 1Gbps, and \fB2\fR for 10Gbps. Valid values range from
4542 1 to 65535.
4543 .RE

4545 .sp
4546 .ne 2
4547 .na
4548 \fB\fBstp_edge\fR\fR

```

```

4549 .ad
4550 .sp .6
4551 .RS 4n
4552 Enables or disables bridge edge port detection. If set to \fB0\fR (false), the
4553 system assumes that the port is connected to other bridges even if no bridge
4554 PDUs of any type are seen. The default value is \fB1\fR, which detects edge
4555 ports automatically.
4556 .RE

4558 .sp
4559 .ne 2
4560 .na
4561 \fB\fBstp_p2p\fR\fR
4562 .ad
4563 .sp .6
4564 .RS 4n
4565 Sets bridge point-to-point operation mode. Possible values are \fBtrue\fR,
4566 \fBfalse\fR, and \fBauto\fR. When set to \fBauto\fR, point-to-point connections
4567 are automatically discovered. When set to \fBtrue\fR, the port mode is forced
4568 to use point-to-point. When set to \fBfalse\fR, the port mode is forced to use
4569 normal multipoint mode. The default value is \fBauto\fR.
4570 .RE

4572 .sp
4573 .ne 2
4574 .na
4575 \fB\fBstp_mcheck\fR\fR
4576 .ad
4577 .sp .6
4578 .RS 4n
4579 Triggers the system to run the RSTP \fBForce BPDU Migration Check\fR procedure
4580 on this link. The procedure is triggered by setting the property value to
4581 \fB1\fR. The property is automatically reset back to \fB0\fR. This value cannot
4582 be set unless the following are true:
4583 .RS +4
4584 .TP
4585 .ie t \(\bu
4586 .el o
4587 The link is bridged
4588 .RE
4589 .RS +4
4590 .TP
4591 .ie t \(\bu
4592 .el o
4593 The bridge is protected by Spanning Tree
4594 .RE
4595 .RS +4
4596 .TP
4597 .ie t \(\bu
4598 .el o
4599 The bridge \fBforce-protocol\fR value is at least 2 (RSTP)
4600 .RE
4601 The default value is 0.
4602 .RE

4604 .sp
4605 .ne 2
4606 .na
4607 \fB\fBzone\fR\fR
4608 .ad
4609 .sp .6
4610 .RS 4n
4611 Specifies the zone to which the link belongs. This property can be modified
4612 only temporarily through \fBdladm\fR, and thus the \fB-t\fR option must be
4613 specified. To modify the zone assignment such that it persists across reboots,
4614 please use \fBzonecfg\fR(1M). Possible values consist of any exclusive-IP zone

```

```

4615 currently running on the system. By default, the zone binding is as per
4616 \fBzonecfg\fR(1M).
4617 .RE

4619 .SS "Wifi Link Properties"
4620 .sp
4621 .LP
4622 The following \fBWiFi\fR link properties are supported. Note that the ability
4623 to set a given property to a given value depends on the driver and hardware.
4624 .sp
4625 .ne 2
4626 .na
4627 \fB\fBchannel\fR\fR
4628 .ad
4629 .sp .6
4630 .RS 4n
4631 Specifies the channel to use. This property can be modified only by certain
4632 \fBWiFi\fR links when in \fBIBSS\fR mode. The default value and allowed range
4633 of values varies by regulatory domain.
4634 .RE

4636 .sp
4637 .ne 2
4638 .na
4639 \fB\fBpowermode\fR\fR
4640 .ad
4641 .sp .6
4642 .RS 4n
4643 Specifies the power management mode of the \fBWiFi\fR link. Possible values are
4644 \fBoff\fR (disable power management), \fBmax\fR (maximum power savings), and
4645 \fBfast\fR (performance-sensitive power management). Default is \fBoff\fR.
4646 .RE

4648 .sp
4649 .ne 2
4650 .na
4651 \fB\fBradio\fR\fR
4652 .ad
4653 .sp .6
4654 .RS 4n
4655 Specifies the radio mode of the \fBWiFi\fR link. Possible values are \fBon\fR
4656 or \fBoff\fR. Default is \fBon\fR.
4657 .RE

4659 .sp
4660 .ne 2
4661 .na
4662 \fB\fBspeed\fR\fR
4663 .ad
4664 .sp .6
4665 .RS 4n
4666 Specifies a fixed speed for the \fBWiFi\fR link, in megabits per second. The
4667 set of possible values depends on the driver and hardware (but is shown by
4668 \fBshow-linkprop\fR); common speeds include 1, 2, 11, and 54. By default, there
4669 is no fixed speed.
4670 .RE

4672 .SS "Ethernet Link Properties"
4673 .sp
4674 .LP
4675 The following MII Properties, as documented in \fBieee802.3\fR(5), are
4676 supported in read-only mode:
4677 .RS +4
4678 .TP
4679 .ie t \bu
4680 .el o

```

```

4681 \fBduplex\fR
4682 .RE
4683 .RS +4
4684 .TP
4685 .ie t \bu
4686 .el o
4687 \fBstate\fR
4688 .RE
4689 .RS +4
4690 .TP
4691 .ie t \bu
4692 .el o
4693 \fBadv_autoneg_cap\fR
4694 .RE
4695 .RS +4
4696 .TP
4697 .ie t \bu
4698 .el o
4699 \fBadv_10fdx_cap\fR
4700 .RE
4701 .RS +4
4702 .TP
4703 .ie t \bu
4704 .el o
4705 \fBadv_1000fdx_cap\fR
4706 .RE
4707 .RS +4
4708 .TP
4709 .ie t \bu
4710 .el o
4711 \fBadv_100hdx_cap\fR
4712 .RE
4713 .RS +4
4714 .TP
4715 .ie t \bu
4716 .el o
4717 \fBadv_100fdx_cap\fR
4718 .RE
4719 .RS +4
4720 .TP
4721 .ie t \bu
4722 .el o
4723 \fBadv_10hdx_cap\fR
4724 .RE
4725 .RS +4
4726 .TP
4727 .ie t \bu
4728 .el o
4729 \fBadv_10fdx_cap\fR
4730 .RE
4731 .RS +4
4732 .TP
4733 .ie t \bu
4734 .el o
4735 \fBadv_10hdx_cap\fR
4736 .RE
4737 .sp
4738 .LP
4739 Each \fBadv\fR property (for example, \fBadv_10fdx_cap\fR) also has a
4740 read/write counterpart \fBen\fR property (for example, \fBen_10fdx_cap\fR)
4741 controlling parameters used at auto-negotiation. In the absence of Power
4742 Management, the \fBadv\fR* speed/duplex parameters provide the values that are
4743 both negotiated and currently effective in hardware. However, with Power
4744 Management enabled, the speed/duplex capabilities currently exposed in hardware
4745 might be a subset of the set of bits that were used in initial link parameter
4746 negotiation. Thus the MII \fBadv\fR parameters are marked read-only, with an

```

```

4747 additional set of \fBen\fR* parameters for configuring speed and duplex
4748 properties at initial negotiation.
4749 .sp
4750 .LP
4751 Note that the \fBadv_autoneg_cap\fR does not have an \fBen_autoneg_cap\fR
4752 counterpart: the \fBadv_autoneg_cap\fR is a 0/1 switch that turns off/on
4753 autonegotiation itself, and therefore cannot be impacted by Power Management.
4754 .sp
4755 .LP
4756 In addition, the following Ethernet properties are reported:
4757 .sp
4758 .ne 2
4759 .na
4760 \fB\fBspeed\fR\fR
4761 .ad
4762 .sp .6
4763 .RS 4n
4764 (read-only) The operating speed of the device, in Mbps.
4765 .RE

4767 .sp
4768 .ne 2
4769 .na
4770 \fB\fBmtu\fR\fR
4771 .ad
4772 .sp .6
4773 .RS 4n
4774 The maximum client SDU (Send Data Unit) supported by the device. Valid range is
4775 68-65536.
4776 .RE

4778 .sp
4779 .ne 2
4780 .na
4781 \fB\fBflowctrl\fR\fR
4782 .ad
4783 .sp .6
4784 .RS 4n
4785 Establishes flow-control modes that will be advertised by the device. Valid
4786 input is one of:
4787 .sp
4788 .ne 2
4789 .na
4790 \fB\fBno\fR\fR
4791 .ad
4792 .sp .6
4793 .RS 4n
4794 No flow control enabled.
4795 .RE

4797 .sp
4798 .ne 2
4799 .na
4800 \fB\fBrx\fR\fR
4801 .ad
4802 .sp .6
4803 .RS 4n
4804 Receive, and act upon incoming pause frames.
4805 .RE

4807 .sp
4808 .ne 2
4809 .na
4810 \fB\fBtx\fR\fR
4811 .ad
4812 .sp .6

```

```

4813 .RS 4n
4814 Transmit pause frames to the peer when congestion occurs, but ignore received
4815 pause frames.
4816 .RE

4818 .sp
4819 .ne 2
4820 .na
4821 \fB\fBbi\fR\fR
4822 .ad
4823 .sp .6
4824 .RS 4n
4825 Bidirectional flow control.
4826 .RE

4828 Note that the actual settings for this value are constrained by the
4829 capabilities allowed by the device and the link partner.
4830 .RE

4832 .sp
4833 .ne 2
4834 .na
4835 \fB\fBtagmode\fR\fR
4836 .ad
4837 .sp .6
4838 .RS 4n
4839 This link property controls the conditions in which 802.1Q VLAN tags will be
4840 inserted in packets being transmitted on the link. Two mode values can be
4841 assigned to this property:
4842 .sp
4843 .ne 2
4844 .na
4845 \fB\fBnormal\fR\fR
4846 .ad
4847 .RS 12n
4848 Insert a VLAN tag in outgoing packets under the following conditions:
4849 .RS +4
4850 .TP
4851 .ie t \(\bu
4852 .el o
4853 The packet belongs to a VLAN.
4854 .RE
4855 .RS +4
4856 .TP
4857 .ie t \(\bu
4858 .el o
4859 The user requested priority tagging.
4860 .RE
4861 .RE

4863 .sp
4864 .ne 2
4865 .na
4866 \fB\fBvlanonly\fR\fR
4867 .ad
4868 .RS 12n
4869 Insert a VLAN tag only when the outgoing packet belongs to a VLAN. If a tag is
4870 being inserted in this mode and the user has also requested a non-zero
4871 priority, the priority is honored and included in the VLAN tag.
4872 .RE

4874 The default value is \fBvlanonly\fR.
4875 .RE

4877 .SS "IP Tunnel Link Properties"
4878 .sp

```

```

4879 .LP
4880 The following IP tunnel link properties are supported.
4881 .sp
4882 .ne 2
4883 .na
4884 \fB\fBhoplimit\fR\fR
4885 .ad
4886 .sp .6
4887 .RS 4n
4888 Specifies the IPv4 TTL or IPv6 hop limit for the encapsulating outer IP header
4889 of a tunnel link. This property exists for all tunnel types. The default value
4890 is 64.
4891 .RE

4893 .sp
4894 .ne 2
4895 .na
4896 \fB\fBencaplimit\fR\fR
4897 .ad
4898 .sp .6
4899 .RS 4n
4900 Specifies the IPv6 encapsulation limit for an IPv6 tunnel as defined in RFC
4901 2473. This value is the tunnel nesting limit for a given tunneled packet. The
4902 default value is 4. A value of 0 disables the encapsulation limit.
4903 .RE

4905 .SH EXAMPLES
4906 .LP
4907 \fBExample 1\fR Configuring an Aggregation
4908 .sp
4909 .LP
4910 To configure a data-link over an aggregation of devices \fBbge0\fR and
4911 \fBbge1\fR with key 1, enter the following command:

4913 .sp
4914 .in +2
4915 .nf
4916 # \fBdladm create-aggr -d bge0 -d bge1 1\fR
4917 .fi
4918 .in -2
4919 .sp

4921 .LP
4922 \fBExample 2\fR Connecting to a WiFi Link
4923 .sp
4924 .LP
4925 To connect to the most optimal available unsecured network on a system with a
4926 single \fBWiFi\fR link (as per the prioritization rules specified for
4927 \fBconnect-wifi\fR), enter the following command:

4929 .sp
4930 .in +2
4931 .nf
4932 # \fBdladm connect-wifi\fR
4933 .fi
4934 .in -2
4935 .sp

4937 .LP
4938 \fBExample 3\fR Creating a WiFi Key
4939 .sp
4940 .LP
4941 To interactively create the \fBWEP\fR key \fBmykey\fR, enter the following
4942 command:
4944 .sp

```

```

4945 .in +2
4946 .nf
4947 # \fBdladm create-secobj -c wep mykey\fR
4948 .fi
4949 .in -2
4950 .sp

4952 .sp
4953 .LP
4954 Alternatively, to non-interactively create the \fBWEP\fR key \fBmykey\fR using
4955 the contents of a file:
4956 .sp
4957 .in +2
4958 .nf
4959 # \fBumask 077\fR
4960 # \fBcat >/tmp/mykey.$$ <<EOF\fR
4961 \fB12345\fR
4962 \fBEOF\fR
4963 # \fBdladm create-secobj -c wep -f /tmp/mykey.$$ mykey\fR
4964 # \fBrm /tmp/mykey.$$>\fR
4965 .fi
4966 .in -2
4967 .sp

4970 .LP
4971 \fBExample 4\fR Connecting to a Specified Encrypted WiFi Link
4972 .sp
4973 .LP
4974 To use key \fBmykey\fR to connect to \fBESSID\fR \fBwlan\fR on link \fBath0\fR,
4975 enter the following command:
4976 .sp
4977 .in +2
4978 .nf
4979 # \fBdladm connect-wifi -k mykey -e wlan ath0\fR
4980 .fi
4981 .in -2
4982 .sp

4985 .LP
4986 \fBExample 5\fR Changing a Link Property
4987 .sp
4988 .LP
4989 To set \fBpowermode\fR to the value \fBfast\fR on link \fBpcwl0\fR, enter the
4990 following command:
4991 .sp
4992 .in +2
4993 .nf
4994 # \fBdladm set-linkprop -p powermode=fast pcwl0\fR
4995 .fi
4996 .in -2
4997 .sp

5000 .LP
5001 \fBExample 6\fR Connecting to a WPA-Protected WiFi Link
5002 .sp
5003 .LP
5004 Create a WPA key \fBpsk\fR and enter the following command:
5005 .sp
5006 .in +2
5007 .nf
5008 # \fBdladm create-secobj -c wpa psk\fR
5009 .fi

```

```

5011 .in -2
5012 .sp

5014 .sp
5015 .LP
5016 To then use key \fBpsk\fR to connect to ESSID \fBwlan\fR on link \fBath0\fR,
5017 enter the following command:
5018 .sp
5019 .in +2
5020 .nf
5021 .fi
5022 # \fBdladm connect-wifi -k psk -e wlan ath0\fR
5023 .fi
5024 .in -2
5025 .sp

5027 .LP
5028 \fBExample 7\fR Renaming a Link
5029 .sp
5030 .LP
5031 To rename the \fBbge0\fR link to \fBmgmt0\fR, enter the following command:
5032 .sp
5033 .in +2
5034 .nf
5035 .fi
5036 # \fBdladm rename-link bge0 mgmt0\fR
5037 .fi
5038 .in -2
5039 .sp

5041 .LP
5042 \fBExample 8\fR Replacing a Network Card
5043 .sp
5044 .LP
5045 Consider that the \fBbge0\fR device, whose link was named \fBmgmt0\fR as shown
5046 in the previous example, needs to be replaced with a \fBce0\fR device because
5047 of a hardware failure. The \fBbge0\fR NIC is physically removed, and replaced
5048 with a new \fBce0\fR NIC. To associate the newly added \fBce0\fR device with
5049 the \fBmgmt0\fR configuration previously associated with \fBbge0\fR, enter the
5050 following command:
5051 .sp
5052 .in +2
5053 .nf
5054 .fi
5055 # \fBdladm rename-link ce0 mgmt0\fR
5056 .fi
5057 .in -2
5058 .sp

5060 .LP
5061 \fBExample 9\fR Removing a Network Card
5062 .sp
5063 .LP
5064 Suppose that in the previous example, the intent is not to replace the
5065 \fBbge0\fR NIC with another NIC, but rather to remove and not replace the
5066 hardware. In that case, the \fBmgmt0\fR datalink configuration is not slated to
5067 be associated with a different physical device as shown in the previous
5068 example, but needs to be deleted. Enter the following command to delete the
5069 datalink configuration associated with the \fBmgmt0\fR datalink, whose physical
5070 hardware (\fBbge0\fR in this case) has been removed:
5071 .sp
5072 .in +2
5073 .nf
5074 .fi
5075 # \fBdladm delete-phys mgmt0\fR
5076 .fi

```

```

5077 .in -2
5078 .sp

5079 .LP
5080 \fBExample 10\fR Using Parseable Output to Capture a Single Field
5081 .sp
5082 .LP
5083 .fi
5084 The following assignment saves the MTU of link \fBnet0\fR to a variable named
5085 \fBmtu\fR.

5086 .sp
5087 .in +2
5088 .nf
5089 .fi
5090 # \fBmtu='dladm show-link -p -o mtu net0'\fR
5091 .fi
5092 .in -2
5093 .sp

5094 .LP
5095 \fBExample 11\fR Using Parseable Output to Iterate over Links
5096 .sp
5097 .LP
5098 .fi
5099 The following script displays the state of each link on the system.

5100 .sp
5101 .in +2
5102 .nf
5103 .fi
5104 # \fBdladm show-link -p -o link,state | while IFS=: read link state; do
5105         print "Link $link is in state $state"
5106     done\fR
5107 .fi
5108 .in -2
5109 .sp

5110 .LP
5111 \fBExample 12\fR Configuring VNICs
5112 .sp
5113 .LP
5114 .fi
5115 Create two VNICs with names \fBhello0\fR and \fBtest1\fR over a single physical
5116 link \fBbge0\fR.

5117 .sp
5118 .in +2
5119 .nf
5120 .fi
5121 # \fBdladm create-vnic -l bge0 hello0\fR
5122 # \fBdladm create-vnic -l bge0 test1\fR
5123 .fi
5124 .in -2
5125 .sp

5126 .LP
5127 \fBExample 13\fR Configuring VNICs and Allocating Bandwidth and Priority
5128 .sp
5129 .LP
5130 .fi
5131 Create two VNICs with names \fBhello0\fR and \fBtest1\fR over a single physical
5132 link \fBbge0\fR and make \fBhello0\fR a high priority VNIC with a
5133 factory-assigned MAC address with a maximum bandwidth of 50 Mbps. Make
5134 \fBtest1\fR a low priority VNIC with a random MAC address and a maximum
5135 bandwidth of 100Mbps.

5136 .sp
5137 .in +2
5138 .nf
5139 .fi
5140 # \fBdladm create-vnic -l bge0 -m factory -p maxbw=50,priority=high hello0\fR
5141 # \fBdladm create-vnic -l bge0 -m random -p maxbw=100M,priority=low test1\fR
5142 .fi

```

```

5143 .in -2
5144 .sp
5146 .LP
5147 \fBExample 14\fR Configuring a VNIC with a Factory MAC Address
5148 .sp
5149 .LP
5150 First, list the available factory MAC addresses and choose one of them:
5152 .sp
5153 .in +2
5154 .nf
5155 # \fBdladm show-phys -m bge0\fR
5156 LINK      SLOT      ADDRESS      INUSE      CLIENT
5157 bge0      primary   0:e0:81:27:d4:47  yes        bge0
5158 bge0      1          8:0:20:fe:4e:a5  no
5159 bge0      2          8:0:20:fe:4e:a6  no
5160 bge0      3          8:0:20:fe:4e:a7  no
5161 .fi
5162 .in -2
5163 .sp
5165 .sp
5166 .LP
5167 Create a VNIC named \fBhello0\fR and use slot 1's address:
5169 .sp
5170 .in +2
5171 .nf
5172 # \fBdladm create-vnic -l bge0 -m factory -n 1 hello0\fR
5173 # \fBdladm show-phys -m bge0\fR
5174 LINK      SLOT      ADDRESS      INUSE      CLIENT
5175 bge0      primary   0:e0:81:27:d4:47  yes        bge0
5176 bge0      1          8:0:20:fe:4e:a5  yes        hello0
5177 bge0      2          8:0:20:fe:4e:a6  no
5178 bge0      3          8:0:20:fe:4e:a7  no
5179 .fi
5180 .in -2
5181 .sp
5183 .LP
5184 \fBExample 15\fR Creating a VNIC with User-Specified MAC Address, Binding it to
5185 Set of Processors
5186 .sp
5187 .LP
5188 Create a VNIC with name \fBhello0\fR, with a user specified MAC address, and a
5189 processor binding \fB0, 1, 2, 3\fR.
5191 .sp
5192 .in +2
5193 .nf
5194 # \fBdladm create-vnic -l bge0 -m 8:0:20:fe:4e:b8 -p cpus=0,1,2,3 hello0\fR
5195 .fi
5196 .in -2
5197 .sp
5199 .LP
5200 \fBExample 16\fR Creating a Virtual Network Without a Physical NIC
5201 .sp
5202 .LP
5203 First, create an etherstub with name \fBstub1\fR:
5205 .sp
5206 .in +2
5207 .nf
5208 # \fBdladm create-etherstub stub1\fR

```

```

5209 .fi
5210 .in -2
5211 .sp
5213 .sp
5214 .LP
5215 Create two VNICs with names \fBhello0\fR and \fBtest1\fR on the etherstub. This
5216 operation implicitly creates a virtual switch connecting \fBhello0\fR and
5217 \fBtest1\fR.
5219 .sp
5220 .in +2
5221 .nf
5222 # \fBdladm create-vnic -l stub1 hello0\fR
5223 # \fBdladm create-vnic -l stub1 test1\fR
5224 .fi
5225 .in -2
5226 .sp
5228 .LP
5229 \fBExample 17\fR Showing Network Usage
5230 .sp
5231 .LP
5232 Network usage statistics can be stored using the extended accounting facility,
5233 \fBacctadm\fR(1M).
5235 .sp
5236 .in +2
5237 .nf
5238 # \fBacctadm -e basic -f /var/log/net.log net\fR
5239 # \fBacctadm net\fR
5240 Network accounting: active
5241 Network accounting file: /var/log/net.log
5242 Tracked Network resources: basic
5243 Untracked Network resources: src_ip,dst_ip,src_port,dst_port,protocol,
5244 dsfield
5245 .fi
5246 .in -2
5247 .sp
5249 .sp
5250 .LP
5251 The saved historical data can be retrieved in summary form using the
5252 \fBshow-usage\fR subcommand:
5254 .sp
5255 .in +2
5256 .nf
5257 # \fBdladm show-usage -f /var/log/net.log\fR
5258 LINK      DURATION  IPACKETS RBYTES  OPACKETS OBYTES      BANDWIDTH
5259 e1000g0    80         1031     546908      0       0      2.44 Kbps
5260 .fi
5261 .in -2
5262 .sp
5264 .LP
5265 \fBExample 18\fR Displaying Bridge Information
5266 .sp
5267 .LP
5268 The following commands use the \fBshow-bridge\fR subcommand with no and various
5269 options.
5271 .sp
5272 .in +2
5273 .nf
5274 # \fBdladm show-bridge\fR

```

```

5275 BRIDGE      PROTECT ADDRESS          PRIORITY DESROOT
5276 foo        stp    32768/8:0:20:bf:f 32768   8192/0:d0:0:76:14:38
5277 bar        stp    32768/8:0:20:e5:8 32768   8192/0:d0:0:76:14:38

5279 # \fBdladm show-bridge -l foo\fR
5280 LINK        STATE     UPTIME   DESROOT
5281 hme0       forwarding 117     8192/0:d0:0:76:14:38
5282 qfe1       forwarding 117     8192/0:d0:0:76:14:38

5284 # \fBdladm show-bridge -s foo\fR
5285 BRIDGE      DROPS    FORWARDS
5286 foo         0        302

5288 # \fBdladm show-bridge -ls foo\fR
5289 LINK        DROPS    RECV     XMIT
5290 hme0       0        360832   31797
5291 qfe1       0        322311   356852

5293 # \fBdladm show-bridge -f foo\fR
5294 DEST        AGE      FLAGS   OUTPUT
5295 8:0:20:bc:a7:dc 10.860  --      hme0
5296 8:0:20:bf:f9:69  --      L       hme0
5297 8:0:20:c0:20:26 17.420  --      hme0
5298 8:0:20:e5:86:11  --      L       qfe1
5299 .fi
5300 .in -2
5301 .sp

5303 .LP
5304 \fBExample 19 \fRCreating an IPv4 Tunnel
5305 .sp
5306 .LP
5307 The following sequence of commands creates and then displays a persistent IPv4
5308 tunnel link named \fBmytunnel0\fR between 66.1.2.3 and 192.4.5.6:

5310 .sp
5311 .in +2
5312 .nf
5313 # \fBdladm create-iptun -T ipv4 -s 66.1.2.3 -d 192.4.5.6 mytunnel0\fR
5314 # \fBdladm show-iptun mytunnel0\fR
5315 LINK        TYPE    FLAGS   SOURCE           DESTINATION
5316 mytunnel0   ipv4   --     66.1.2.3       192.4.5.6
5317 .fi
5318 .in -2
5319 .sp

5321 .sp
5322 .LP
5323 A point-to-point IP interface can then be created over this tunnel link:

5325 .sp
5326 .in +2
5327 .nf
5328 # \fBifconfig mytunnel0 plumb 10.1.0.1 10.1.0.2 up\fR
5329 .fi
5330 .in -2
5331 .sp

5333 .sp
5334 .LP
5335 As with any other IP interface, configuration persistence for this IP interface
5336 is achieved by placing the desired \fBifconfig\fR commands (in this case, the
5337 command for "\fB10.1.0.1 10.1.0.2\fR") into \fB/etc/hostname.mytunnel0\fR.

5339 .LP
5340 \fBExample 20 \fRCreating a 6to4 Tunnel

```

```

5341 .sp
5342 .LP
5343 The following command creates a 6to4 tunnel link. The IPv4 address of the 6to4
5344 router is 75.10.11.12.

5346 .sp
5347 .in +2
5348 .nf
5349 # \fBdladm create-iptun -T 6to4 -s 75.10.11.12 sitetunnel0\fR
5350 # \fBdladm show-iptun sitetunnel0\fR
5351 LINK        TYPE    FLAGS   SOURCE           DESTINATION
5352 sitetunnel0  6to4   --     75.10.11.12      --
5353 .fi
5354 .in -2
5355 .sp

5357 .sp
5358 .LP
5359 The following command plumbs an IPv6 interface on this tunnel:
5361 .sp
5362 .in +2
5363 .nf
5364 # \fBifconfig sitetunnel0 inet6 plumb up\fR
5365 # \fBifconfig sitetunnel0 inet6\fR
5366 sitetunnel0: flags=2200041 <UP,RUNNING,NONUD,IPv6> mtu 65515 index 3
5367     inet tunnel src 75.10.11.12
5368     tunnel hop limit 64
5369     inet6 2002:4b0a:b0c::1/16
5370 .fi
5371 .in -2
5372 .sp

5374 .sp
5375 .LP
5376 Note that the system automatically configures the IPv6 address on the 6to4 IP
5377 interface. See \fBifconfig\fR(1M) for a description of how IPv6 addresses are
5378 configured on 6to4 tunnel links.

5380 .SH ATTRIBUTES
5381 .sp
5382 .LP
5383 See \fBAttributes\fR(5) for descriptions of the following attributes:
5384 .sp
5385 .LP
5386 \fB/usr/sbin\fR
5387 .sp

5389 .sp
5390 .TS
5391 box;
5392 c | c
5393 l | l .
5394 ATTRIBUTE TYPE   ATTRIBUTE VALUE
5395 Interface Stability     Committed
5396 .TE

5399 .sp
5400 .LP
5401 \fB/sbin\fR
5402 .sp

5404 .sp
5405 .TS
5406 box;

```

```
5407 c | c
5408 1 | 1 .
5409 ATTRIBUTE TYPE ATTRIBUTE VALUE
5410 -
5411 Interface Stability      Committed
5412 .TE

5414 .SH SEE ALSO
5415 .sp
5416 .LP
5417 \fBacctadm\fR(1M), \fBautopush\fR(1M), \fBifconfig\fR(1M), \fBipsecconf\fR(1M),
5418 \fBndd\fR(1M), \fBpsrset\fR(1M), \fBwpad\fR(1M), \fBzonecfg\fR(1M),
5419 \fBattributes\fR(5), \fBieee802.3\fR(5), \fBdlpi\fR(7P)
5420 .SH NOTES
5421 .sp
5422 .LP
5423 The preferred method of referring to an aggregation in the aggregation
5424 subcommands is by its link name. Referring to an aggregation by its integer
5425 \fIkey\fR is supported for backward compatibility, but is not necessary. When
5426 creating an aggregation, if a \fIkey\fR is specified instead of a link name,
5427 the aggregation's link name will be automatically generated by \fBdladm\fR as
5428 \fBaggr\fR\fIkey\fR.
```

\*\*\*\*\*  
213717 Thu Feb 20 18:59:05 2014  
new/usr/src/uts/common/io/mac/mac.c  
2553 mac address should be a dladm link property  
\*\*\*\*\*  
\_\_\_\_\_ unchanged\_portion\_omitted \_\_\_\_\_

```
2846 /*
2847 * Checks the size of the value size specified for a property as
2848 * part of a property operation. Returns B_TRUE if the size is
2849 * correct, B_FALSE otherwise.
2850 */
2851 boolean_t
2852 mac_prop_check_size(mac_prop_id_t id, uint_t valsiz, boolean_t is_range)
2853 {
2854     uint_t minsize = 0;
2855
2856     if (is_range)
2857         return (valsiz >= sizeof (mac_propval_range_t));
2858
2859     switch (id) {
2860     case MAC_PROP_ZONE:
2861         minsize = sizeof (dld_ioc_zid_t);
2862         break;
2863     case MAC_PROP_AUTOUPUSH:
2864         if (valsiz != 0)
2865             minsize = sizeof (struct dlautopush);
2866         break;
2867     case MAC_PROP_TAGMODE:
2868         minsize = sizeof (link_tagmode_t);
2869         break;
2870     case MAC_PROP_RESOURCE:
2871     case MAC_PROP_RESOURCE_EFF:
2872         minsize = sizeof (mac_resource_props_t);
2873         break;
2874     case MAC_PROP_DUPLEX:
2875         minsize = sizeof (link_duplex_t);
2876         break;
2877     case MAC_PROP_SPEED:
2878         minsize = sizeof (uint64_t);
2879         break;
2880     case MAC_PROP_STATUS:
2881         minsize = sizeof (link_state_t);
2882         break;
2883     case MAC_PROP_AUTONEG:
2884     case MAC_PROP_EN_AUTONEG:
2885         minsize = sizeof (uint8_t);
2886         break;
2887     case MAC_PROP_MTU:
2888     case MAC_PROP_LLIMIT:
2889     case MAC_PROP_LDECAY:
2890         minsize = sizeof (uint32_t);
2891         break;
2892     case MAC_PROP_FLOWCTRL:
2893         minsize = sizeof (link_flowctrl_t);
2894         break;
2895     case MAC_PROP_ADV_10GFDX_CAP:
2896     case MAC_PROP_EN_10GFDX_CAP:
2897     case MAC_PROP_ADV_1000HDX_CAP:
2898     case MAC_PROP_EN_1000HDX_CAP:
2899     case MAC_PROP_ADV_100FDX_CAP:
2900     case MAC_PROP_EN_100FDX_CAP:
2901     case MAC_PROP_ADV_100HDX_CAP:
2902     case MAC_PROP_EN_100HDX_CAP:
2903     case MAC_PROP_ADV_10FDX_CAP:
2904     case MAC_PROP_EN_10FDX_CAP:
```

```
2905     case MAC_PROP_ADV_10HDX_CAP:
2906     case MAC_PROP_EN_10HDX_CAP:
2907     case MAC_PROP_ADV_100T4_CAP:
2908     case MAC_PROP_EN_100T4_CAP:
2909         minsize = sizeof (uint8_t);
2910         break;
2911     case MAC_PROP_PVID:
2912         minsize = sizeof (uint16_t);
2913         break;
2914     case MAC_PROP_IPTUN_HOPLIMIT:
2915         minsize = sizeof (uint32_t);
2916         break;
2917     case MAC_PROP_IPTUN_ENCAPLIMIT:
2918         minsize = sizeof (uint32_t);
2919         break;
2920     case MAC_PROP_MAX_TX_RINGS_AVAIL:
2921     case MAC_PROP_MAX_RX_RINGS_AVAIL:
2922     case MAC_PROP_MAX_RXHWCLNT_AVAIL:
2923     case MAC_PROP_MAX_TXHWCLNT_AVAIL:
2924         minsize = sizeof (uint_t);
2925         break;
2926     case MAC_PROP_WL_ESSID:
2927         minsize = sizeof (wl_linkstatus_t);
2928         break;
2929     case MAC_PROP_WL_BSSID:
2930         minsize = sizeof (wl_bssid_t);
2931         break;
2932     case MAC_PROP_WL_BSSTYPE:
2933         minsize = sizeof (wl_bss_type_t);
2934         break;
2935     case MAC_PROP_WL_LINKSTATUS:
2936         minsize = sizeof (wl_linkstatus_t);
2937         break;
2938     case MAC_PROP_WL_DESIRED_RATES:
2939         minsize = sizeof (wl_rates_t);
2940         break;
2941     case MAC_PROP_WL_SUPPORTED_RATES:
2942         minsize = sizeof (wl_rates_t);
2943         break;
2944     case MAC_PROP_WL_AUTH_MODE:
2945         minsize = sizeof (wl_authmode_t);
2946         break;
2947     case MAC_PROP_WL_ENCRYPTION:
2948         minsize = sizeof (wl_encryption_t);
2949         break;
2950     case MAC_PROP_WL_RSSI:
2951         minsize = sizeof (wl_rssi_t);
2952         break;
2953     case MAC_PROP_WL_PHY_CONFIG:
2954         minsize = sizeof (wl_phy_conf_t);
2955         break;
2956     case MAC_PROP_WL_CAPABILITY:
2957         minsize = sizeof (wl_capability_t);
2958         break;
2959     case MAC_PROP_WL_WPA:
2960         minsize = sizeof (wl_wpa_t);
2961         break;
2962     case MAC_PROP_WL_SCANRESULTS:
2963         minsize = sizeof (wl_wpa_ess_t);
2964         break;
2965     case MAC_PROP_WL_POWER_MODE:
2966         minsize = sizeof (wl_ps_mode_t);
2967         break;
2968     case MAC_PROP_WL_RADIO:
2969         minsize = sizeof (wl_radio_t);
2970         break;
```

```

2971     case MAC_PROP_WL_ESS_LIST:
2972         minsize = sizeof (wl_ess_list_t);
2973         break;
2974     case MAC_PROP_WL_KEY_TAB:
2975         minsize = sizeof (wl_wep_key_tab_t);
2976         break;
2977     case MAC_PROP_WL_CREATE_IBSS:
2978         minsize = sizeof (wl_create_ibss_t);
2979         break;
2980     case MAC_PROP_WL_SETOPTIE:
2981         minsize = sizeof (wl_wpa_ie_t);
2982         break;
2983     case MAC_PROP_WL_DELKEY:
2984         minsize = sizeof (wl_del_key_t);
2985         break;
2986     case MAC_PROP_WL_KEY:
2987         minsize = sizeof (wl_key_t);
2988         break;
2989     case MAC_PROP_WL_MLME:
2990         minsize = sizeof (wl_mlme_t);
2991         break;
2992     case MAC_PROP_MACADDRESS:
2993         minsize = sizeof (mac_addrprop_t);
2994 #endif /* ! codereview */
2995     }
2996
2997     return (valsize >= minsize);
2998 }
2999
3000 /*
3001 * mac_set_prop() sets MAC or hardware driver properties:
3002 *
3003 * - MAC-managed properties such as resource properties include maxbw,
3004 * priority, and cpu binding list, as well as the default port VID
3005 * used by bridging. These properties are consumed by the MAC layer
3006 * itself and not passed down to the driver. For resource control
3007 * properties, this function invokes mac_set_resources() which will
3008 * cache the property value in mac_impl_t and may call
3009 * mac_client_set_resource() to update property value of the primary
3010 * mac client, if it exists.
3011 *
3012 * - Properties which act on the hardware and must be passed to the
3013 * driver, such as MTU, through the driver's mc_setprop() entry point.
3014 */
3015 int
3016 mac_set_prop(mac_handle_t mh, mac_prop_id_t id, char *name, void *val,
3017   uint_t valsize)
3018 {
3019     int err = ENOTSUP;
3020     mac_impl_t *mip = (mac_impl_t *)mh;
3021
3022     ASSERT(MAC_PERIM_HELD(mh));
3023
3024     switch (id) {
3025     case MAC_PROP_RESOURCE: {
3026         mac_resource_props_t *mrp;
3027
3028         /* call mac_set_resources() for MAC properties */
3029         ASSERT(valsize >= sizeof (mac_resource_props_t));
3030         mrp = kmem_zalloc(sizeof (*mrp), KM_SLEEP);
3031         bcopy(val, mrp, sizeof (*mrp));
3032         err = mac_set_resources(mh, mrp);
3033         kmem_free(mrp, sizeof (*mrp));
3034         break;
3035     }
3036
3037     case MAC_PROP_PVID:
3038         ASSERT(valsize >= sizeof (uint16_t));
3039         if (mip->mi_state_flags & MIS_IS_VNIC)
3040             return (EINVAL);
3041         err = mac_set_pvid(mh, *(uint16_t *)val);
3042         break;
3043
3044     case MAC_PROP_MTU: {
3045         uint32_t mtu;
3046
3047         ASSERT(valsize >= sizeof (uint32_t));
3048         bcopy(val, &mtu, sizeof (mtu));
3049         err = mac_set_mtu(mh, mtu, NULL);
3050         break;
3051     }
3052
3053     case MAC_PROP_LLIMIT:
3054     case MAC_PROP_LDECAY: {
3055         uint32_t learnval;
3056
3057         if (valsize < sizeof (learnval) ||
3058             (mip->mi_state_flags & MIS_IS_VNIC))
3059             return (EINVAL);
3060         bcopy(val, &learnval, sizeof (learnval));
3061         if (learnval == 0 && id == MAC_PROP_LDECAY)
3062             return (EINVAL);
3063         if (id == MAC_PROP_LLIMIT)
3064             mip->mi_llimit = learnval;
3065         else
3066             mip->mi_ldecay = learnval;
3067         err = 0;
3068         break;
3069     }
3070
3071     case MAC_PROP_MACADDRESS: {
3072         mac_addrprop_t *addrprop = val;
3073
3074         if (addrprop->ma_len != mip->mi_type->mt_addr_length)
3075             return (EINVAL);
3076
3077         err = mac_unicast_primary_set(mh, addrprop->ma_addr);
3078         break;
3079     }
3080
3081 #endif /* ! codereview */
3082     default:
3083         /* For other driver properties, call driver's callback */
3084         if (mip->mi_callbacks->mc_callbacks & MC_SetProp) {
3085             err = mip->mi_callbacks->mc_SetProp(mip->mi_driver,
3086                                                 name, id, valsize, val);
3087         }
3088     }
3089     return (err);
3090 }
3091
3092 /*
3093 * mac_get_prop() gets MAC or device driver properties.
3094 *
3095 * If the property is a driver property, mac_get_prop() calls driver's callback
3096 * entry point to get it.
3097 * If the property is a MAC property, mac_get_prop() invokes mac_get_resources()
3098 * which returns the cached value in mac_impl_t.
3099 */
3100 int
3101 mac_get_prop(mac_handle_t mh, mac_prop_id_t id, char *name, void *val,
3102   uint_t valsize)
3103 }
```

```

3037     case MAC_PROP_PVID:
3038         ASSERT(valsize >= sizeof (uint16_t));
3039         if (mip->mi_state_flags & MIS_IS_VNIC)
3040             return (EINVAL);
3041         err = mac_set_pvid(mh, *(uint16_t *)val);
3042         break;
3043
3044     case MAC_PROP_MTU: {
3045         uint32_t mtu;
3046
3047         ASSERT(valsize >= sizeof (uint32_t));
3048         bcopy(val, &mtu, sizeof (mtu));
3049         err = mac_set_mtu(mh, mtu, NULL);
3050         break;
3051     }
3052
3053     case MAC_PROP_LLIMIT:
3054     case MAC_PROP_LDECAY: {
3055         uint32_t learnval;
3056
3057         if (valsize < sizeof (learnval) ||
3058             (mip->mi_state_flags & MIS_IS_VNIC))
3059             return (EINVAL);
3060         bcopy(val, &learnval, sizeof (learnval));
3061         if (learnval == 0 && id == MAC_PROP_LDECAY)
3062             return (EINVAL);
3063         if (id == MAC_PROP_LLIMIT)
3064             mip->mi_llimit = learnval;
3065         else
3066             mip->mi_ldecay = learnval;
3067         err = 0;
3068         break;
3069     }
3070
3071     case MAC_PROP_MACADDRESS: {
3072         mac_addrprop_t *addrprop = val;
3073
3074         if (addrprop->ma_len != mip->mi_type->mt_addr_length)
3075             return (EINVAL);
3076
3077         err = mac_unicast_primary_set(mh, addrprop->ma_addr);
3078         break;
3079     }
3080
3081 #endif /* ! codereview */
3082     default:
3083         /* For other driver properties, call driver's callback */
3084         if (mip->mi_callbacks->mc_callbacks & MC_SetProp) {
3085             err = mip->mi_callbacks->mc_SetProp(mip->mi_driver,
3086                                                 name, id, valsize, val);
3087         }
3088     }
3089     return (err);
3090 }
3091
3092 /*
3093 * mac_get_prop() gets MAC or device driver properties.
3094 *
3095 * If the property is a driver property, mac_get_prop() calls driver's callback
3096 * entry point to get it.
3097 * If the property is a MAC property, mac_get_prop() invokes mac_get_resources()
3098 * which returns the cached value in mac_impl_t.
3099 */
3100 int
3101 mac_get_prop(mac_handle_t mh, mac_prop_id_t id, char *name, void *val,
3102   uint_t valsize)
3103 }
```

```

3103 {
3104     int err = ENOTSUP;
3105     mac_impl_t *mip = (mac_impl_t *)mh;
3106     uint_t rings;
3107     uint_t vlinks;
3108
3109     bzero(val, valsiz);
3110
3111     switch (id) {
3112         case MAC_PROP_RESOURCE: {
3113             mac_resource_props_t *mrp;
3114
3115             /* If mac property, read from cache */
3116             ASSERT(valsiz >= sizeof (mac_resource_props_t));
3117             mrp = kmem_zalloc(sizeof (*mrp), KM_SLEEP);
3118             mac_get_resources(mh, mrp);
3119             bcopy(mrp, val, sizeof (*mrp));
3120             kmem_free(mrp, sizeof (*mrp));
3121             return (0);
3122         }
3123         case MAC_PROP_RESOURCE_EFF: {
3124             mac_resource_props_t *mrp;
3125
3126             /* If mac effective property, read from client */
3127             ASSERT(valsiz >= sizeof (mac_resource_props_t));
3128             mrp = kmem_zalloc(sizeof (*mrp), KM_SLEEP);
3129             mac_get_effective_resources(mh, mrp);
3130             bcopy(mrp, val, sizeof (*mrp));
3131             kmem_free(mrp, sizeof (*mrp));
3132             return (0);
3133         }
3134
3135         case MAC_PROP_PVID:
3136             ASSERT(valsiz >= sizeof (uint16_t));
3137             if (mip->mi_state_flags & MIS_IS_VNIC)
3138                 return (EINVAL);
3139             *(uint16_t *)val = mac_get_pvid(mh);
3140             return (0);
3141
3142         case MAC_PROP_LLIMIT:
3143         case MAC_PROP_LDECAY:
3144             ASSERT(valsiz >= sizeof (uint32_t));
3145             if (mip->mi_state_flags & MIS_IS_VNIC)
3146                 return (EINVAL);
3147             if (id == MAC_PROP_LLIMIT)
3148                 bcopy(&mip->mi_llimit, val, sizeof (mip->mi_llimit));
3149             else
3150                 bcopy(&mip->mi_ldecay, val, sizeof (mip->mi_ldecay));
3151             return (0);
3152
3153         case MAC_PROP_MTU: {
3154             uint32_t sdu;
3155
3156             ASSERT(valsiz >= sizeof (uint32_t));
3157             mac_sdu_get2(mh, NULL, &sdu, NULL);
3158             bcopy(&sdu, val, sizeof (sdu));
3159
3160             return (0);
3161         }
3162         case MAC_PROP_STATUS: {
3163             link_state_t link_state;
3164
3165             if (valsiz < sizeof (link_state))
3166                 return (EINVAL);
3167             link_state = mac_link_get(mh);
3168             bcopy(&link_state, val, sizeof (link_state));
3169         }
3170     }
3171     return (0);
3172 }
3173
3174 case MAC_PROP_MAX_RX_RINGS_AVAIL:
3175 case MAC_PROP_MAX_TX_RINGS_AVAIL:
3176     ASSERT(valsiz >= sizeof (uint_t));
3177     rings = id == MAC_PROP_MAX_RX_RINGS_AVAIL ?
3178             mac_rxavail_get(mh) : mac_txavail_get(mh);
3179     bcopy(&rings, val, sizeof (uint_t));
3180     return (0);
3181
3182 case MAC_PROP_MAX_RXHWCLNT_AVAIL:
3183 case MAC_PROP_MAX_TXHWCLNT_AVAIL:
3184     ASSERT(valsiz >= sizeof (uint_t));
3185     vlinks = id == MAC_PROP_MAX_RXHWCLNT_AVAIL ?
3186             mac_rxhwlnksavail_get(mh) : mac_txhwlnksavail_get(mh);
3187     bcopy(&vlinks, val, sizeof (uint_t));
3188     return (0);
3189
3190 case MAC_PROP_RXRINGSRANGE:
3191 case MAC_PROP_TXRINGSRANGE:
3192     /*
3193      * The value for these properties are returned through
3194      * the MAC_PROP_RESOURCE property.
3195      */
3196     return (0);
3197
3198 case MAC_PROP_MACADDRESS: {
3199     mac_addrprop_t *addrprop = val;
3200
3201     if (valsiz < sizeof (mac_addrprop_t))
3202         return (EINVAL);
3203     mac_unicast_primary_get(mh, addrprop->ma_addr);
3204     addrprop->ma_len = mip->mi_type->mt_addr_length;
3205     return (0);
3206 }
3207 #endif /* ! codereview */
3208 default:
3209     break;
3210 }
3211 }
3212
3213 /* If driver property, request from driver */
3214 if (mip->mi_callbacks->mc_callbacks & MC_GETPROP) {
3215     err = mip->mi_callbacks->mc_getprop(mip->mi_driver, name, id,
3216                                         valsiz, val);
3217 }
3218
3219 return (err);
3220 }
3221
3222 /*
3223  * Helper function to initialize the range structure for use in
3224  * mac_get_prop. If the type can be other than uint32, we can
3225  * pass that as an arg.
3226  */
3227 static void
3228 _mac_set_range(mac_propval_range_t *range, uint32_t min, uint32_t max)
3229 {
3230     range->mpr_count = 1;
3231     range->mpr_type = MAC_PROPVAL_UINT32;
3232     range->mpr_range_uint32[0].mpur_min = min;
3233     range->mpr_range_uint32[0].mpur_max = max;
3234 }
```

```

3170
3171     return (0);
3172 }
3173
3174 case MAC_PROP_MAX_RX_RINGS_AVAIL:
3175 case MAC_PROP_MAX_TX_RINGS_AVAIL:
3176     ASSERT(valsiz >= sizeof (uint_t));
3177     rings = id == MAC_PROP_MAX_RX_RINGS_AVAIL ?
3178             mac_rxavail_get(mh) : mac_txavail_get(mh);
3179     bcopy(&rings, val, sizeof (uint_t));
3180     return (0);
3181
3182 case MAC_PROP_MAX_RXHWCLNT_AVAIL:
3183 case MAC_PROP_MAX_TXHWCLNT_AVAIL:
3184     ASSERT(valsiz >= sizeof (uint_t));
3185     vlinks = id == MAC_PROP_MAX_RXHWCLNT_AVAIL ?
3186             mac_rxhwlnksavail_get(mh) : mac_txhwlnksavail_get(mh);
3187     bcopy(&vlinks, val, sizeof (uint_t));
3188     return (0);
3189
3190 case MAC_PROP_RXRINGSRANGE:
3191 case MAC_PROP_TXRINGSRANGE:
3192     /*
3193      * The value for these properties are returned through
3194      * the MAC_PROP_RESOURCE property.
3195      */
3196     return (0);
3197
3198 case MAC_PROP_MACADDRESS: {
3199     mac_addrprop_t *addrprop = val;
3200
3201     if (valsiz < sizeof (mac_addrprop_t))
3202         return (EINVAL);
3203     mac_unicast_primary_get(mh, addrprop->ma_addr);
3204     addrprop->ma_len = mip->mi_type->mt_addr_length;
3205     return (0);
3206 }
3207 #endif /* ! codereview */
3208 default:
3209     break;
3210 }
3211 }
3212
3213 /* If driver property, request from driver */
3214 if (mip->mi_callbacks->mc_callbacks & MC_GETPROP) {
3215     err = mip->mi_callbacks->mc_getprop(mip->mi_driver, name, id,
3216                                         valsiz, val);
3217 }
3218
3219 return (err);
3220 }
3221
3222 /*
3223  * Helper function to initialize the range structure for use in
3224  * mac_get_prop. If the type can be other than uint32, we can
3225  * pass that as an arg.
3226  */
3227 static void
3228 _mac_set_range(mac_propval_range_t *range, uint32_t min, uint32_t max)
3229 {
3230     range->mpr_count = 1;
3231     range->mpr_type = MAC_PROPVAL_UINT32;
3232     range->mpr_range_uint32[0].mpur_min = min;
3233     range->mpr_range_uint32[0].mpur_max = max;
3234 }
```

```

3236 /*
3237  * Returns information about the specified property, such as default
3238  * values or permissions.
3239 */
3240 int
3241 mac_prop_info(mac_handle_t mh, mac_prop_id_t id, char *name,
3242  void *default_val, uint_t default_size, mac_propval_range_t *range,
3243  uint_t *perm)
3244 {
3245     mac_prop_info_state_t state;
3246     mac_impl_t *mip = (mac_impl_t *)mh;
3247     uint_t max;
3248
3249     /*
3250      * A property is read/write by default unless the driver says
3251      * otherwise.
3252     */
3253     if (perm != NULL)
3254         *perm = MAC_PROP_PERM_RW;
3255
3256     if (default_val != NULL)
3257         bzero(default_val, default_size);
3258
3259     /*
3260      * First, handle framework properties for which we don't need to
3261      * involve the driver.
3262     */
3263     switch (id) {
3264     case MAC_PROP_RESOURCE:
3265     case MAC_PROP_PVID:
3266     case MAC_PROP_LLIMIT:
3267     case MAC_PROP_LDECAY:
3268         return (0);
3269
3270     case MAC_PROP_MAX_RX_RINGS_AVAIL:
3271     case MAC_PROP_MAX_TX_RINGS_AVAIL:
3272     case MAC_PROP_MAX_RXHWCLNT_AVAIL:
3273     case MAC_PROP_MAX_TXHWCLNT_AVAIL:
3274         if (perm != NULL)
3275             *perm = MAC_PROP_PERM_READ;
3276         return (0);
3277
3278     case MAC_PROP_RXRINGS RANGE:
3279     case MAC_PROP_TXRINGS RANGE:
3280         /*
3281          * Currently, we support range for RX and TX rings properties.
3282          * When we extend this support to maxbw, cpus and priority,
3283          * we should move this to mac_get_resources.
3284          * There is no default value for RX or TX rings.
3285        */
3286     if ((mip->mi_state_flags & MIS_IS_VNIC) &&
3287         mac_is_vnic_primary(mh)) {
3288         /*
3289          * We don't support setting rings for a VLAN
3290          * data link because it shares its ring with the
3291          * primary MAC client.
3292        */
3293     if (perm != NULL)
3294         *perm = MAC_PROP_PERM_READ;
3295     if (range != NULL)
3296         range->mpr_count = 0;
3297     } else if (range != NULL) {
3298         if (mip->mi_state_flags & MIS_IS_VNIC)
3299             mh = mac_get_lower_mac_handle(mh);
3300         mip = (mac_impl_t *)mh;

```

```

3301
3302     if ((id == MAC_PROP_RXRINGS RANGE &&
3303         mip->mi_rx_group_type == MAC_GROUP_TYPE_STATIC) ||
3304         (id == MAC_PROP_TXRINGS RANGE &&
3305         mip->mi_tx_group_type == MAC_GROUP_TYPE_STATIC)) {
3306         if (id == MAC_PROP_RXRINGS RANGE) {
3307             if ((mac_rxhwlnksavail_get(mh) +
3308                 mac_rxhwlnksrsvd_get(mh)) <= 1) {
3309                 /*
3310                  * doesn't support groups or
3311                  * rings
3312                */
3313             range->mpr_count = 0;
3314         } else {
3315             /*
3316              * supports specifying groups,
3317              * but not rings
3318            */
3319             _mac_set_range(range, 0, 0);
3320         }
3321     } else {
3322         if ((mac_txhwlnksavail_get(mh) +
3323             mac_txhwlnksrsvd_get(mh)) <= 1) {
3324             /*
3325               * doesn't support groups or
3326               * rings
3327             */
3328             range->mpr_count = 0;
3329         } else {
3330             /*
3331               * supports specifying groups,
3332               * but not rings
3333             */
3334             _mac_set_range(range, 0, 0);
3335         }
3336     }
3337     max = id == MAC_PROP_RXRINGS RANGE ?
3338         mac_rxavail_get(mh) + mac_rxrsvd_get(mh) :
3339         mac_txavail_get(mh) + mac_txrsvd_get(mh);
3340     if (max <= 1) {
3341         /*
3342          * doesn't support groups or
3343          * rings
3344        */
3345     } else {
3346         range->mpr_count = 0;
3347     }
3348     /*
3349      * -1 because we have to leave out the
3350      * default ring.
3351    */
3352     _mac_set_range(range, 1, max - 1);
3353 }
3354 }
3355 }
3356 }
3357 }
3358 }
3359 }
3360 }
3361 }
3362 }
3363 }
3364 }
3365 }
3366 }

case MAC_PROP_STATUS:
if (perm != NULL)
    *perm = MAC_PROP_PERM_READ;
return (0);

case MAC_PROP_MACADDRESS: {
    mac_addrprop_t *defaddr = default_val;
    if (defaddr != NULL) {
        if (default_size < sizeof (mac_addrprop_t))

```

```

3367         return (EINVAL);
3368         bcopy(mip->mi_info.mi_unicst_addr, defaddr->ma_addr,
3369               mip->mi_type->mt_addr_length);
3370         defaddr->ma_len = mip->mi_type->mt_addr_length;
3371     }
3372     return (0);
3373 }
3374 #endif /* ! codereview */
3375 }

3377 /*
3378 * Get the property info from the driver if it implements the
3379 * property info entry point.
3380 */
3381 bzero(&state, sizeof (state));

3383 if (mip->mi_callbacks->mc_callbacks & MC_PROPINFO) {
3384     state.pr_default = default_val;
3385     state.pr_default_size = default_size;

3387 /*
3388 * The caller specifies the maximum number of ranges
3389 * it can accomodate using mpr_count. We don't touch
3390 * this value until the driver returns from its
3391 * mc_propinfo() callback, and ensure we don't exceed
3392 * this number of range as the driver defines
3393 * supported range from its mc_propinfo().
3394 *
3395 * pr_range_cur_count keeps track of how many ranges
3396 * were defined by the driver from its mc_propinfo()
3397 * entry point.
3398 *
3399 * On exit, the user-specified range mpr_count returns
3400 * the number of ranges specified by the driver on
3401 * success, or the number of ranges it wanted to
3402 * define if that number of ranges could not be
3403 * accomodated by the specified range structure. In
3404 * the latter case, the caller will be able to
3405 * allocate a larger range structure, and query the
3406 * property again.
3407 */
3408 state.pr_range_cur_count = 0;
3409 state.pr_range = range;

3411 mip->mi_callbacks->mc_propinfo(mip->mi_driver, name, id,
3412     (mac_prop_info_handle_t)&state);

3414 if (state.pr_flags & MAC_PROP_INFO_RANGE)
3415     range->mpr_count = state.pr_range_cur_count;

3417 /*
3418 * The operation could fail if the buffer supplied by
3419 * the user was too small for the range or default
3420 * value of the property.
3421 */
3422 if (state.pr_errno != 0)
3423     return (state.pr_errno);

3425 if (perm != NULL && state.pr_flags & MAC_PROP_INFO_PERM)
3426     *perm = state.pr_perm;
3427 }

3429 /*
3430 * The MAC layer may want to provide default values or allowed
3431 * ranges for properties if the driver does not provide a
3432 * property info entry point, or that entry point exists, but

```

```

3433         * it did not provide a default value or allowed ranges for
3434         * that property.
3435         */
3436     switch (id) {
3437     case MAC_PROP_MTU: {
3438         uint32_t sdu;

3440         mac_sdu_get2(mh, NULL, &sdu, NULL);

3442         if (range != NULL && !(state.pr_flags &
3443             MAC_PROP_INFO_RANGE)) {
3444             /* MTU range */
3445             _mac_set_range(range, sdu, sdu);
3446         }

3448         if (default_val != NULL && !(state.pr_flags &
3449             MAC_PROP_INFO_DEFAULT)) {
3450             if (mip->mi_info.mi_media == DL_ETHER)
3451                 sdu = ETHERMTU;
3452             /* default MTU value */
3453             bcopy(&sdu, default_val, sizeof (sdu));
3454         }
3455     }
3456 }

3458     return (0);
3459 }

3461 int
3462 mac_fastpath_disable(mac_handle_t mh)
3463 {
3464     mac_impl_t *mip = (mac_impl_t *)mh;
3466     if ((mip->mi_state_flags & MIS_LEGACY) == 0)
3467         return (0);
3469     return (mip->mi_capab_legacy.ml_fastpath_disable(mip->mi_driver));
3470 }

3472 void
3473 mac_fastpath_enable(mac_handle_t mh)
3474 {
3475     mac_impl_t *mip = (mac_impl_t *)mh;
3477     if ((mip->mi_state_flags & MIS_LEGACY) == 0)
3478         return;
3480     mip->mi_capab_legacy.ml_fastpath_enable(mip->mi_driver);
3481 }

3483 void
3484 mac_register_priv_prop(mac_impl_t *mip, char **priv_props)
3485 {
3486     uint_t nprops, i;
3488     if (priv_props == NULL)
3489         return;
3491     nprops = 0;
3492     while (priv_props[nprops] != NULL)
3493         nprops++;
3494     if (nprops == 0)
3495         return;
3498     mip->mi_priv_prop = kmalloc(nprops * sizeof (char *), KM_SLEEP);

```

```

3500     for (i = 0; i < nprops; i++) {
3501         mip->mi_priv_prop[i] = kmalloc(MAXLINKPROPNAME, KM_SLEEP);
3502         (void) strlcpy(mip->mi_priv_prop[i], priv_props[i],
3503                         MAXLINKPROPNAME);
3504     }
3506     mip->mi_priv_prop_count = nprops;
3507 }
3509 void
3510 mac_unregister_priv_prop(mac_impl_t *mip)
3511 {
3512     uint_t i;
3514     if (mip->mi_priv_prop_count == 0) {
3515         ASSERT(mip->mi_priv_prop == NULL);
3516         return;
3517     }
3519     for (i = 0; i < mip->mi_priv_prop_count; i++)
3520         kmem_free(mip->mi_priv_prop[i], MAXLINKPROPNAME);
3521     kmem_free(mip->mi_priv_prop, mip->mi_priv_prop_count *
3522               sizeof (char *));
3524     mip->mi_priv_prop = NULL;
3525     mip->mi_priv_prop_count = 0;
3526 }
3528 /*
3529 * mac_ring_t 'mr' macros. Some rogue drivers may access ring structure
3530 * (by invoking mac_rx()) even after processing mac_stop_ring(). In such
3531 * cases if MAC free's the ring structure after mac_stop_ring(), any
3532 * illegal access to the ring structure coming from the driver will panic
3533 * the system. In order to protect the system from such inadvertent access,
3534 * we maintain a cache of rings in the mac_impl_t after they get free'd up.
3535 * When packets are received on free'd up rings, MAC (through the generation
3536 * count mechanism) will drop such packets.
3537 */
3538 static mac_ring_t *
3539 mac_ring_alloc(mac_impl_t *mip)
3540 {
3541     mac_ring_t *ring;
3543     mutex_enter(&mip->mi_ring_lock);
3544     if (mip->mi_ring_freelist != NULL) {
3545         ring = mip->mi_ring_freelist;
3546         mip->mi_ring_freelist = ring->mr_next;
3547         bzero(ring, sizeof (mac_ring_t));
3548         mutex_exit(&mip->mi_ring_lock);
3549     } else {
3550         mutex_exit(&mip->mi_ring_lock);
3551         ring = kmem_cache_alloc(mac_ring_cache, KM_SLEEP);
3552     }
3553     ASSERT((ring != NULL) && (ring->mr_state == MR_FREE));
3554     return (ring);
3555 }
3557 static void
3558 mac_ring_free(mac_impl_t *mip, mac_ring_t *ring)
3559 {
3560     ASSERT(ring->mr_state == MR_FREE);
3562     mutex_enter(&mip->mi_ring_lock);
3563     ring->mr_state = MR_FREE;
3564     ring->mr_flag = 0;

```

```

3565     ring->mr_next = mip->mi_ring_freelist;
3566     ring->mr_mip = NULL;
3567     mip->mi_ring_freelist = ring;
3568     mac_ring_stat_delete(ring);
3569     mutex_exit(&mip->mi_ring_lock);
3570 }
3572 static void
3573 mac_ring_freeall(mac_impl_t *mip)
3574 {
3575     mac_ring_t *ring_next;
3576     mutex_enter(&mip->mi_ring_lock);
3577     mac_ring_t *ring = mip->mi_ring_freelist;
3578     while (ring != NULL) {
3579         ring_next = ring->mr_next;
3580         kmem_cache_free(mac_ring_cache, ring);
3581         ring = ring_next;
3582     }
3583     mip->mi_ring_freelist = NULL;
3584     mutex_exit(&mip->mi_ring_lock);
3585 }
3587 int
3588 mac_start_ring(mac_ring_t *ring)
3589 {
3590     int rv = 0;
3592     ASSERT(ring->mr_state == MR_FREE);
3594     if (ring->mr_start != NULL) {
3595         rv = ring->mr_start(ring->mr_driver, ring->mr_gen_num);
3596         if (rv != 0)
3597             return (rv);
3598     }
3600     ring->mr_state = MR_INUSE;
3601     return (rv);
3602 }
3604 void
3605 mac_stop_ring(mac_ring_t *ring)
3606 {
3607     ASSERT(ring->mr_state == MR_INUSE);
3609     if (ring->mr_stop != NULL)
3610         ring->mr_stop(ring->mr_driver);
3612     ring->mr_state = MR_FREE;
3614     /*
3615      * Increment the ring generation number for this ring.
3616      */
3617     ring->mr_gen_num++;
3618 }
3620 int
3621 mac_start_group(mac_group_t *group)
3622 {
3623     int rv = 0;
3625     if (group->mrg_start != NULL)
3626         rv = group->mrg_start(group->mrg_driver);
3628     return (rv);
3629 }

```

```

3631 void
3632 mac_stop_group(mac_group_t *group)
3633 {
3634     if (group->mrg_stop != NULL)
3635         group->mrg_stop(group->mrg_driver);
3636 }
3638 /*
3639  * Called from mac_start() on the default Rx group. Broadcast and multicast
3640  * packets are received only on the default group. Hence the default group
3641  * needs to be up even if the primary client is not up, for the other groups
3642  * to be functional. We do this by calling this function at mac_start time
3643  * itself. However the broadcast packets that are received can't make their
3644  * way beyond mac_rx until a mac client creates a broadcast flow.
3645 */
3646 static int
3647 mac_start_group_and_rings(mac_group_t *group)
3648 {
3649     mac_ring_t      *ring;
3650     int             rv = 0;
3652
3653     ASSERT(group->mrg_state == MAC_GROUP_STATE_REGISTERED);
3654     if ((rv = mac_start_group(group)) != 0)
3655         return (rv);
3656
3657     for (ring = group->mrg_rings; ring != NULL; ring = ring->mr_next) {
3658         ASSERT(ring->mr_state == MR_FREE);
3659         if ((rv = mac_start_ring(ring)) != 0)
3660             goto error;
3661         ring->mr_classify_type = MAC_SW_CLASSIFIER;
3662     }
3663     return (0);
3664
3665 error:
3666     mac_stop_group_and_rings(group);
3667     return (rv);
3668 }
3669 /* Called from mac_stop on the default Rx group */
3670 static void
3671 mac_stop_group_and_rings(mac_group_t *group)
3672 {
3673     mac_ring_t      *ring;
3674
3675     for (ring = group->mrg_rings; ring != NULL; ring = ring->mr_next) {
3676         if (ring->mr_state != MR_FREE) {
3677             mac_stop_ring(ring);
3678             ring->mr_flag = 0;
3679             ring->mr_classify_type = MAC_NO_CLASSIFIER;
3680         }
3681     }
3682     mac_stop_group(group);
3683 }
3684
3685 static mac_ring_t *
3686 mac_init_ring(mac_impl_t *mip, mac_group_t *group, int index,
3687                 mac_capab_rings_t *cap_rings)
3688 {
3689     mac_ring_t      *ring, *rnext;
3690     mac_ring_info_t ring_info;
3691     ddi_intr_handle_t ddi_handle;
3692
3693     ring = mac_ring_alloc(mip);
3694
3695     /* Prepare basic information of ring */

```

```

3698
3699     /*
3700      * Ring index is numbered to be unique across a particular device.
3701      * Ring index computation makes following assumptions:
3702      *   - For drivers with static grouping (e.g. ixgbe, bge),
3703      *     ring index exchanged with the driver (e.g. during mr_rget)
3704      *     is unique only across the group the ring belongs to.
3705      *   - Drivers with dynamic grouping (e.g. nxge), start
3706      *     with single group (mrg_index = 0).
3707
3708     ring->mr_index = group->mrg_index * group->mrg_info.mgi_count + index;
3709     ring->mr_type = group->mrg_type;
3710     ring->mrg_gn = (mac_group_handle_t)group;
3711
3712     /* Insert the new ring to the list. */
3713     ring->mr_next = group->mrg_rings;
3714     group->mrg_rings = ring;
3715
3716     /* Zero to reuse the info data structure */
3717     bzero(&ring_info, sizeof (ring_info));
3718
3719     /* Query ring information from driver */
3720     cap_rings->mr_rget(mip->mi_driver, group->mrg_type, group->mrg_index,
3721                         index, &ring_info, (mac_ring_handle_t)ring);
3722
3723     ring->mr_info = ring_info;
3724
3725     /*
3726      * The interrupt handle could be shared among multiple rings.
3727      * Thus if there is a bunch of rings that are sharing an
3728      * interrupt, then only one ring among the bunch will be made
3729      * available for interrupt re-targeting; the rest will have
3730      * ddi_shared flag set to TRUE and would not be available for
3731      * be interrupt re-targeting.
3732
3733     if ((ddi_handle = ring_info.mri_intr.mi_ddi_handle) != NULL) {
3734         rnext = ring->mr_next;
3735         while (rnext != NULL) {
3736             if (rnext->mr_info.mri_intr.mi_ddi_handle ==
3737                 ddi_handle) {
3738                 /*
3739                  * If default ring (mr_index == 0) is part
3740                  * of a group of rings sharing an
3741                  * interrupt, then set ddi_shared flag for
3742                  * the default ring and give another ring
3743                  * the chance to be re-targeted.
3744
3745                 if (rnext->mr_index == 0 &&
3746                     !rnext->mr_info.mri_intr.mi_ddi_shared) {
3747                     rnext->mr_info.mri_intr.mi_ddi_shared =
3748                         B_TRUE;
3749                 } else {
3750                     ring->mr_info.mri_intr.mi_ddi_shared =
3751                         B_TRUE;
3752                 }
3753             }
3754             rnext = rnext->mr_next;
3755         }
3756
3757         /*
3758          * If rnext is NULL, then no matching ddi_handle was found.
3759          * Rx rings get registered first. So if this is a Tx ring,
3760          * then go through all the Rx rings and see if there is a
3761          * matching ddi handle.
3762
3763         if (rnext == NULL && ring->mr_type == MAC_RING_TYPE_TX) {

```

```

3763             mac_compare_ddi_handle(mip->mi_rx_groups,
3764                                     mip->mi_rx_group_count, ring);
3765     }
3766 }
3767
3768 /* Update ring's status */
3769 ring->mr_state = MR_FREE;
3770 ring->mr_flag = 0;
3771
3772 /* Update the ring count of the group */
3773 group->mrg_cur_count++;
3774
3775 /* Create per ring kstats */
3776 if (ring->mr_stat != NULL) {
3777     ring->mr_mip = mip;
3778     mac_ring_stat_create(ring);
3779 }
3780
3781 return (ring);
3782 }
3783
3784 /*
3785 * Rings are chained together for easy regrouping.
3786 */
3787 static void
3788 mac_init_group(mac_impl_t *mip, mac_group_t *group, int size,
3789                 mac_capab_rings_t *cap_rings)
3790 {
3791     int index;
3792
3793     /*
3794      * Initialize all ring members of this group. Size of zero will not
3795      * enter the loop, so it's safe for initializing an empty group.
3796      */
3797     for (index = size - 1; index >= 0; index--)
3798         (void) mac_init_ring(mip, group, index, cap_rings);
3799 }
3800
3801 int
3802 mac_init_rings(mac_impl_t *mip, mac_ring_type_t rtype)
3803 {
3804     mac_capab_rings_t      *cap_rings;
3805     mac_group_t            *group;
3806     mac_group_t            *groups;
3807     mac_group_info_t       *group_info;
3808     uint_t                  group_free = 0;
3809     uint_t                  ring_left;
3810     mac_ring_t              *ring;
3811     int                     g;
3812     int                     err = 0;
3813     uint_t                  grpCnt;
3814     boolean_t               pseudo_txgrp = B_FALSE;
3815
3816     switch (rtype) {
3817     case MAC_RING_TYPE_RX:
3818         ASSERT(mip->mi_rx_groups == NULL);
3819
3820         cap_rings = &mip->mi_rx_rings_cap;
3821         cap_rings->mr_type = MAC_RING_TYPE_RX;
3822         break;
3823     case MAC_RING_TYPE_TX:
3824         ASSERT(mip->mi_tx_groups == NULL);
3825
3826         cap_rings = &mip->mi_tx_rings_cap;
3827         cap_rings->mr_type = MAC_RING_TYPE_TX;
3828         break;

```

```

3829     default:
3830         ASSERT(B_FALSE);
3831     }
3832
3833     if (!i_mac_capab_get((mac_handle_t)mip, MAC_CAPAB_RINGS, cap_rings))
3834         return (0);
3835     grpCnt = cap_rings->mr_gnum;
3836
3837     /*
3838      * If we have multiple TX rings, but only one TX group, we can
3839      * create pseudo TX groups (one per TX ring) in the MAC layer,
3840      * except for an agr. For an agr currently we maintain only
3841      * one group with all the rings (for all its ports), going
3842      * forwards we might change this.
3843      */
3844     if (rtype == MAC_RING_TYPE_RX &&
3845         cap_rings->mr_gnum == 0 && cap_rings->mr_rnum > 0 &&
3846         (mip->mi_state_flags & MIS_IS_AGGR) == 0) {
3847         /*
3848          * The -1 here is because we create a default TX group
3849          * with all the rings in it.
3850          */
3851         grpCnt = cap_rings->mr_rnum - 1;
3852         pseudo_txgrp = B_TRUE;
3853     }
3854
3855     /*
3856      * Allocate a contiguous buffer for all groups.
3857      */
3858     groups = kmem_zalloc(sizeof (mac_group_t) * (grpCnt+ 1), KM_SLEEP);
3859
3860     ring_left = cap_rings->mr_rnum;
3861
3862     /*
3863      * Get all ring groups if any, and get their ring members
3864      * if any.
3865      */
3866     for (g = 0; g < grpCnt; g++) {
3867         group = groups + g;
3868
3869         /*
3870          * Prepare basic information of the group */
3871         group->mrg_index = g;
3872         group->mrg_type = rtype;
3873         group->mrg_state = MAC_GROUP_STATE_UNINIT;
3874         group->mrg_mh = (mac_handle_t)mip;
3875         group->mrg_next = group + 1;
3876
3877         /*
3878          * Zero to reuse the info data structure */
3879         bzero(&group_info, sizeof (group_info));
3880
3881         if (pseudo_txgrp) {
3882             /*
3883              * This is a pseudo group that we created, apart
3884              * from setting the state there is nothing to be
3885              * done.
3886              */
3887             group->mrg_state = MAC_GROUP_STATE_REGISTERED;
3888             group_free++;
3889             continue;
3890         }
3891         /*
3892          * Query group information from driver */
3893         cap_rings->mr_gget(mip->mi_driver, rtype, g, &group_info,
3894                             (mac_group_handle_t)group);
3895
3896         switch (cap_rings->mr_group_type) {
3897             case MAC_GROUP_TYPE_DYNAMIC:

```

```

3895     if (cap_rings->mr_gaddring == NULL ||  
3896         cap_rings->mr_gremring == NULL) {  
3897         DTRACE_PROBE3(  
3898             mac_init_rings_no_addremring,  
3899             char *, mip->mi_name,  
3900             mac_group_add_ring_t,  
3901             cap_rings->mr_gaddring,  
3902             mac_group_add_ring_t,  
3903             cap_rings->mr_gremring);  
3904         err = EINVAL;  
3905         goto bail;  
3906     }  
3907  
3908     switch (rtype) {  
3909         case MAC_RING_TYPE_RX:  
3910             /*  
3911                 * The first RX group must have non-zero  
3912                 * rings, and the following groups must  
3913                 * have zero rings.  
3914             */  
3915             if (g == 0 && group_info.mgi_count == 0) {  
3916                 DTRACE_PROBE1(  
3917                     mac_init_rings_rx_def_zero,  
3918                     char *, mip->mi_name);  
3919                 err = EINVAL;  
3920                 goto bail;  
3921             }  
3922             if (g > 0 && group_info.mgi_count != 0) {  
3923                 DTRACE_PROBE3(  
3924                     mac_init_rings_rx_nonzero,  
3925                     char *, mip->mi_name,  
3926                     int, g, int, group_info.mgi_count);  
3927                     err = EINVAL;  
3928                     goto bail;  
3929             }  
3930             break;  
3931         case MAC_RING_TYPE_TX:  
3932             /*  
3933                 * All TX ring groups must have zero rings.  
3934             */  
3935             if (group_info.mgi_count != 0) {  
3936                 DTRACE_PROBE3(  
3937                     mac_init_rings_tx_nonzero,  
3938                     char *, mip->mi_name,  
3939                     int, g, int, group_info.mgi_count);  
3940                     err = EINVAL;  
3941                     goto bail;  
3942             }  
3943             break;  
3944         }  
3945         break;  
3946     case MAC_GROUP_TYPE_STATIC:  
3947         /*  
3948             * Note that an empty group is allowed, e.g., an aggr  
3949             * would start with an empty group.  
3950         */  
3951         break;  
3952     default:  
3953         /* unknown group type */  
3954         DTRACE_PROBE2(mac_init_rings_unknown_type,  
3955             char *, mip->mi_name,  
3956             int, cap_rings->mr_group_type);  
3957         err = EINVAL;  
3958         goto bail;  
3959     }

```

```

3962     /*  
3963         * Driver must register group->mgi_addmac/remmac() for rx groups  
3964         * to support multiple MAC addresses.  
3965     */  
3966     if (rtype == MAC_RING_TYPE_RX) {  
3967         if ((group_info.mgi_addmac == NULL) ||  
3968             (group_info.mgi_addmac == NULL)) {  
3969             goto bail;  
3970         }  
3971     }  
3972  
3973     /* Cache driver-supplied information */  
3974     group->mrg_info = group_info;  
3975  
3976     /* Update the group's status and group count. */  
3977     mac_set_group_state(group, MAC_GROUP_STATE_REGISTERED);  
3978     group_free++;  
3979  
3980     group->mrg_rings = NULL;  
3981     group->mrg_cur_count = 0;  
3982     mac_init_group(mip, group, group_info.mgi_count, cap_rings);  
3983     ring_left -= group_info.mgi_count;  
3984  
3985     /* The current group size should be equal to default value */  
3986     ASSERT(group->mrg_cur_count == group_info.mgi_count);  
3987  
3988     /* Build up a dummy group for free resources as a pool */  
3989     group = groups + grpcnt;  
3990  
3991     /* Prepare basic information of the group */  
3992     group->mrg_index = -1;  
3993     group->mrg_type = rtype;  
3994     group->mrg_state = MAC_GROUP_STATE_UNINIT;  
3995     group->mrg_mh = (mac_handle_t)mip;  
3996     group->mrg_next = NULL;  
3997  
3998     /*  
3999         * If there are ungrouped rings, allocate a continuous buffer for  
4000         * remaining resources.  
4001     */  
4002     if (ring_left != 0) {  
4003         group->mrg_rings = NULL;  
4004         group->mrg_cur_count = 0;  
4005         mac_init_group(mip, group, ring_left, cap_rings);  
4006  
4007         /* The current group size should be equal to ring_left */  
4008         ASSERT(group->mrg_cur_count == ring_left);  
4009  
4010         ring_left = 0;  
4011  
4012         /* Update this group's status */  
4013         mac_set_group_state(group, MAC_GROUP_STATE_REGISTERED);  
4014     } else  
4015         group->mrg_rings = NULL;  
4016  
4017     ASSERT(ring_left == 0);  
4018  
4019     4020 bail:  
4021  
4022     /* Cache other important information to finalize the initialization */  
4023     switch (rtype) {  
4024         case MAC_RING_TYPE_RX:  
4025             mip->mi_rx_group_type = cap_rings->mr_group_type;  
4026             mip->mi_rx_group_count = cap_rings->mr_gnum;

```

```

4027         mip->mi_rx_groups = groups;
4028         mip->mi_rx_donor_grp = groups;
4029         if (mip->mi_rx_group_type == MAC_GROUP_TYPE_DYNAMIC) {
4030             /*
4031                 * The default ring is reserved since it is
4032                 * used for sending the broadcast etc. packets.
4033                 */
4034             mip->mi_rxrings_avail =
4035                 mip->mi_rx_groups->mrg_cur_count - 1;
4036             mip->mi_rxrings_rsvd = 1;
4037         }
4038         /*
4039             * The default group cannot be reserved. It is used by
4040             * all the clients that do not have an exclusive group.
4041             */
4042         mip->mi_rxhwclnt_avail = mip->mi_rx_group_count - 1;
4043         mip->mi_rxhwclnt_used = 1;
4044         break;
4045     case MAC_RING_TYPE_TX:
4046         mip->mi_tx_group_type = pseudo_txgrp ? MAC_GROUP_TYPE_DYNAMIC :
4047             cap_rings->mr_group_type;
4048         mip->mi_tx_group_count = grpCnt;
4049         mip->mi_tx_group_free = group_free;
4050         mip->mi_tx_groups = groups;
4051
4052         group = groups + grpCnt;
4053         ring = group->mrg_rings;
4054         /*
4055             * The ring can be NULL in the case of aggr. Aggr will
4056             * have an empty Tx group which will get populated
4057             * later when pseudo Tx rings are added after
4058             * mac_register() is done.
4059             */
4060         if (ring == NULL) {
4061             ASSERT(mip->mi_state_flags & MIS_IS_AGGR);
4062             /*
4063                 * pass the group to aggr so it can add Tx
4064                 * rings to the group later.
4065                 */
4066             cap_rings->mr_gget(mip->mi_driver, rtype, 0, NULL,
4067                 (mac_group_handle_t)group);
4068             /*
4069                 * Even though there are no rings at this time
4070                 * (rings will come later), set the group
4071                 * state to registered.
4072                 */
4073             group->mrg_state = MAC_GROUP_STATE_REGISTERED;
4074         } else {
4075             /*
4076                 * Ring 0 is used as the default one and it could be
4077                 * assigned to a client as well.
4078                 */
4079             while ((ring->mr_index != 0) && (ring->mr_next != NULL))
4080                 ring = ring->mr_next;
4081             ASSERT(ring->mr_index == 0);
4082             mip->mi_default_tx_ring = (mac_ring_handle_t)ring;
4083         }
4084         if (mip->mi_tx_group_type == MAC_GROUP_TYPE_DYNAMIC)
4085             mip->mi_txrings_avail = group->mrg_cur_count - 1;
4086             /*
4087                 * The default ring cannot be reserved.
4088                 */
4089             mip->mi_txrings_rsvd = 1;
4090
4091         /*
4092             * The default group cannot be reserved. It will be shared
4093             * by clients that do not have an exclusive group.
4094
4095             */
4096             break;
4097         default:
4098             ASSERT(B_FALSE);
4099         }
4100         if (err != 0)
4101             mac_free_rings(mip, rtype);
4102
4103         return (err);
4104
4105     /*
4106         * The ddi interrupt handle could be shared among rings. If so, compare
4107         * the new ring's ddi handle with the existing ones and set ddi_shared
4108         * flag.
4109         */
4110     void
4111     mac_compare_ddi_handle(mac_group_t *groups, uint_t grpCnt, mac_ring_t *cring)
4112     {
4113         mac_group_t *group;
4114         mac_ring_t *ring;
4115         ddi_intr_handle_t ddi_handle;
4116         int g;
4117
4118         ddi_handle = cring->mr_info.mri_intr.mi_ddi_handle;
4119         for (g = 0; g < grpCnt; g++) {
4120             group = groups + g;
4121             for (ring = group->mrg_rings; ring != NULL;
4122                 ring = ring->mr_next) {
4123                 if (ring == cring)
4124                     continue;
4125                 if (ring->mr_info.mri_intr.mi_ddi_handle ==
4126                     ddi_handle) {
4127                     if (cring->mr_type == MAC_RING_TYPE_RX &&
4128                         ring->mr_index == 0 &&
4129                         !ring->mr_info.mri_intr.mi_ddi_shared) {
4130                         ring->mr_info.mri_intr.mi_ddi_shared =
4131                             B_TRUE;
4132                     } else {
4133                         cring->mr_info.mri_intr.mi_ddi_shared =
4134                             B_TRUE;
4135                     }
4136                 }
4137             }
4138         }
4139     }
4140
4141     }
4142 }
4143
4144 /*
4145     * Called to free all groups of particular type (RX or TX). It's assumed that
4146     * no clients are using these groups.
4147     */
4148 void
4149 mac_free_rings(mac_impl_t *mip, mac_ring_type_t rtype)
4150 {
4151     mac_group_t *group, *groups;
4152     uint_t groupCount;
4153
4154     switch (rtype) {
4155     case MAC_RING_TYPE_RX:
4156         if (mip->mi_rx_groups == NULL)
4157             return;
4158
4159         /*
4160             */
4161         break;
4162     case MAC_RING_TYPE_TX:
4163         /*
4164             */
4165         break;
4166     }
4167 }
```

```

4168         /*
4169             */
4170         break;
4171     }
4172
4173     /*
4174         */
4175     break;
4176 }
4177
4178 /*
4179     */
4180 void
4181 mac_free_rings(mac_impl_t *mip, mac_ring_type_t rtype)
4182 {
4183     mac_group_t *group, *groups;
4184     uint_t groupCount;
4185
4186     switch (rtype) {
4187     case MAC_RING_TYPE_RX:
4188         if (mip->mi_rx_groups == NULL)
4189             return;
4190
4191         /*
4192             */
4193         break;
4194     case MAC_RING_TYPE_TX:
4195         /*
4196             */
4197         break;
4198     }
4199 }
```

```

4159     groups = mip->mi_rx_groups;
4160     group_count = mip->mi_rx_group_count;
4161
4162     mip->mi_rx_groups = NULL;
4163     mip->mi_rx_donor_grp = NULL;
4164     mip->mi_rx_group_count = 0;
4165     break;
4166 case MAC_RING_TYPE_TX:
4167     ASSERT(mip->mi_tx_group_count == mip->mi_tx_group_free);
4168
4169     if (mip->mi_tx_groups == NULL)
4170         return;
4171
4172     groups = mip->mi_tx_groups;
4173     group_count = mip->mi_tx_group_count;
4174
4175     mip->mi_tx_groups = NULL;
4176     mip->mi_tx_group_count = 0;
4177     mip->mi_tx_group_free = 0;
4178     mip->mi_default_tx_ring = NULL;
4179     break;
4180 default:
4181     ASSERT(B_FALSE);
4182 }
4183
4184 for (group = groups; group != NULL; group = group->mrg_next) {
4185     mac_ring_t *ring;
4186
4187     if (group->mrg_cur_count == 0)
4188         continue;
4189
4190     ASSERT(group->mrg_rings != NULL);
4191
4192     while ((ring = group->mrg_rings) != NULL) {
4193         group->mrg_rings = ring->mr_next;
4194         mac_ring_free(mip, ring);
4195     }
4196 }
4197
4198 /* Free all the cached rings */
4199 mac_ring_freeall(mip);
4200 /* Free the block of group data strutures */
4201 kmem_free(groups, sizeof (mac_group_t) * (group_count + 1));
4202 }
4203
4204 /*
4205 * Associate a MAC address with a receive group.
4206 *
4207 * The return value of this function should always be checked properly, because
4208 * any type of failure could cause unexpected results. A group can be added
4209 * or removed with a MAC address only after it has been reserved. Ideally,
4210 * a successful reservation always leads to calling mac_group_addmac() to
4211 * steer desired traffic. Failure of adding an unicast MAC address doesn't
4212 * always imply that the group is functioning abnormally.
4213 *
4214 * Currently this function is called everywhere, and it reflects assumptions
4215 * about MAC addresses in the implementation. CR 6735196.
4216 */
4217 int
4218 mac_group_addmac(mac_group_t *group, const uint8_t *addr)
4219 {
4220     ASSERT(group->mrg_type == MAC_RING_TYPE_RX);
4221     ASSERT(group->mrg_info.mgi_addmac != NULL);
4222
4223     return (group->mrg_info.mgi_addmac(group->mrg_info.mgi_driver, addr));
4224 }
```

```

4225     /*
4226      * Remove the association between MAC address and receive group.
4227      */
4228     int
4229     mac_group_remmac(mac_group_t *group, const uint8_t *addr)
4230 {
4231     ASSERT(group->mrg_type == MAC_RING_TYPE_RX);
4232     ASSERT(group->mrg_info.mgi_remmac != NULL);
4233
4234     return (group->mrg_info.mgi_remmac(group->mrg_info.mgi_driver, addr));
4235 }
4236
4237 /*
4238  * This is the entry point for packets transmitted through the bridging code.
4239  * If no bridge is in place, MAC_RING_TX transmits using tx ring. The 'rh'
4240  * pointer may be NULL to select the default ring.
4241  */
4242 mblk_t *
4243 mac_bridge_tx(mac_impl_t *mip, mac_ring_handle_t rh, mblk_t *mp)
4244 {
4245     mac_handle_t mh;
4246
4247     /*
4248      * Once we take a reference on the bridge link, the bridge
4249      * module itself can't unload, so the callback pointers are
4250      * stable.
4251      */
4252     mutex_enter(&mip->mi_bridge_lock);
4253     if ((mh = mip->mi_bridge_link) != NULL)
4254         mac_bridge_ref_cb(mh, B_TRUE);
4255     mutex_exit(&mip->mi_bridge_lock);
4256     if (mh == NULL) {
4257         MAC_RING_TX(mip, rh, mp, mp);
4258     } else {
4259         mp = mac_bridge_tx_cb(mh, rh, mp);
4260         mac_bridge_ref_cb(mh, B_FALSE);
4261     }
4262
4263     return (mp);
4264 }
4265
4266 /*
4267  * Find a ring from its index.
4268  */
4269 mac_ring_handle_t
4270 mac_find_ring(mac_group_handle_t gh, int index)
4271 {
4272     mac_group_t *group = (mac_group_t *)gh;
4273     mac_ring_t *ring = group->mrg_rings;
4274
4275     for (ring = group->mrg_rings; ring != NULL; ring = ring->mr_next)
4276         if (ring->mr_index == index)
4277             break;
4278
4279     return ((mac_ring_handle_t)ring);
4280 }
4281
4282 /*
4283  * Add a ring to an existing group.
4284  */
4285
4286  * The ring must be either passed directly (for example if the ring
4287  * movement is initiated by the framework), or specified through a driver
4288  * index (for example when the ring is added by the driver.
4289  */
4290
4291  * The caller needs to call mac_perim_enter() before calling this function.
4292 */
```

```

4291 int
4292 i_mac_group_add_ring(mac_group_t *group, mac_ring_t *ring, int index)
4293 {
4294     mac_impl_t *mip = (mac_impl_t *)group->mrg_mh;
4295     mac_capab_rings_t *cap_rings;
4296     boolean_t driver_call = (ring == NULL);
4297     mac_group_type_t group_type;
4298     int ret = 0;
4299     flow_entry_t *flent;

4301     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));

4303     switch (group->mrg_type) {
4304         case MAC_RING_TYPE_RX:
4305             cap_rings = &mip->mi_rx_rings_cap;
4306             group_type = mip->mi_rx_group_type;
4307             break;
4308         case MAC_RING_TYPE_TX:
4309             cap_rings = &mip->mi_tx_rings_cap;
4310             group_type = mip->mi_tx_group_type;
4311             break;
4312         default:
4313             ASSERT(B_FALSE);
4314     }

4316     /*
4317      * There should be no ring with the same ring index in the target
4318      * group.
4319     */
4320     ASSERT(mac_find_ring((mac_group_handle_t)group,
4321         driver_call ? index : ring->mr_index) == NULL);

4323     if (driver_call) {
4324         /*
4325          * The function is called as a result of a request from
4326          * a driver to add a ring to an existing group, for example
4327          * from the aggregation driver. Allocate a new mac_ring_t
4328          * for that ring.
4329         */
4330         ring = mac_init_ring(mip, group, index, cap_rings);
4331         ASSERT(group->mrg_state > MAC_GROUP_STATE_UNINIT);
4332     } else {
4333         /*
4334          * The function is called as a result of a MAC layer request
4335          * to add a ring to an existing group. In this case the
4336          * ring is being moved between groups, which requires
4337          * the underlying driver to support dynamic grouping,
4338          * and the mac_ring_t already exists.
4339         */
4340         ASSERT(group_type == MAC_GROUP_TYPE_DYNAMIC);
4341         ASSERT(group->mrg_driver == NULL ||
4342             cap_rings->mr_gaddring != NULL);
4343         ASSERT(ring->mr_gh == NULL);
4344     }

4346     /*
4347      * At this point the ring should not be in use, and it should be
4348      * off the right for the target group.
4349     */
4350     ASSERT(ring->mr_state < MR_INUSE);
4351     ASSERT(ring->mr_srs == NULL);
4352     ASSERT(ring->mr_type == group->mrg_type);

4354     if (!driver_call) {
4355         /*
4356          * Add the driver level hardware ring if the process was not

```

```

4357             * initiated by the driver, and the target group is not the
4358             * group.
4359             */
4360             if (group->mrg_driver != NULL) {
4361                 cap_rings->mr_gaddring(group->mrg_driver,
4362                     ring->mr_driver, ring->mr_type);
4363             }
4364
4365             /*
4366              * Insert the ring ahead existing rings.
4367             */
4368             ring->mr_next = group->mrg_rings;
4369             group->mrg_rings = ring;
4370             ring->mr_gh = (mac_group_handle_t)group;
4371             group->mrg_cur_count++;
4372         }

4374         /*
4375          * If the group has not been actively used, we're done.
4376         */
4377         if (group->mrg_index != -1 &&
4378             group->mrg_state < MAC_GROUP_STATE_RESERVED)
4379             return (0);

4381         /*
4382          * Start the ring if needed. Failure causes to undo the grouping action.
4383         */
4384         if (ring->mr_state != MR_INUSE) {
4385             if ((ret = mac_start_ring(ring)) != 0) {
4386                 if (!driver_call) {
4387                     cap_rings->mr_gremring(group->mrg_driver,
4388                         ring->mr_driver, ring->mr_type);
4389                 }
4390                 group->mrg_cur_count--;
4391                 group->mrg_rings = ring->mr_next;
4392
4393                 ring->mr_gh = NULL;
4394
4395                 if (driver_call)
4396                     mac_ring_free(mip, ring);
4397
4398                 return (ret);
4399             }
4400         }

4402         /*
4403          * Set up SRS/SR according to the ring type.
4404         */
4405         switch (ring->mr_type) {
4406             case MAC_RING_TYPE_RX:
4407                 /*
4408                  * Setup SRS on top of the new ring if the group is
4409                  * reserved for someones exclusive use.
4410                 */
4411                 if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4412                     mac_client_impl_t *mcip;
4413
4414                     mcip = MAC_GROUP_ONLY_CLIENT(group);
4415
4416                     /*
4417                      * Even though this group is reserved we might still
4418                      * have multiple clients, i.e a VLAN shares the
4419                      * group with the primary mac client.
4420                     */
4421                     if (mcip != NULL) {
4422                         flent = mcip->mci_flent;
4423                         ASSERT(flent->fe_rx_srs_cnt > 0);
4424
4425                     }
4426                 }
4427             }
4428         }
4429     }
4430
4431     /*
4432      * Set up SRS/SR according to the ring type.
4433     */
4434     switch (ring->mr_type) {
4435         case MAC_RING_TYPE_RX:
4436             /*
4437               * Setup SRS on top of the new ring if the group is
4438               * reserved for someones exclusive use.
4439             */
4440             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4441                 mac_client_impl_t *mcip;
4442
4443                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4444
4445                 /*
4446                   * Even though this group is reserved we might still
4447                   * have multiple clients, i.e a VLAN shares the
4448                   * group with the primary mac client.
4449                   */
4450                 if (mcip != NULL) {
4451                     flent = mcip->mci_flent;
4452                     ASSERT(flent->fe_rx_srs_cnt > 0);
4453
4454                 }
4455             }
4456         }
4457     }
4458
4459     /*
4460      * Set up SRS/SR according to the ring type.
4461     */
4462     switch (ring->mr_type) {
4463         case MAC_RING_TYPE_RX:
4464             /*
4465               * Setup SRS on top of the new ring if the group is
4466               * reserved for someones exclusive use.
4467             */
4468             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4469                 mac_client_impl_t *mcip;
4470
4471                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4472
4473                 /*
4474                   * Even though this group is reserved we might still
4475                   * have multiple clients, i.e a VLAN shares the
4476                   * group with the primary mac client.
4477                   */
4478                 if (mcip != NULL) {
4479                     flent = mcip->mci_flent;
4480                     ASSERT(flent->fe_rx_srs_cnt > 0);
4481
4482                 }
4483             }
4484         }
4485     }
4486
4487     /*
4488      * Set up SRS/SR according to the ring type.
4489     */
4490     switch (ring->mr_type) {
4491         case MAC_RING_TYPE_RX:
4492             /*
4493               * Setup SRS on top of the new ring if the group is
4494               * reserved for someones exclusive use.
4495             */
4496             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4497                 mac_client_impl_t *mcip;
4498
4499                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4500
4501                 /*
4502                   * Even though this group is reserved we might still
4503                   * have multiple clients, i.e a VLAN shares the
4504                   * group with the primary mac client.
4505                   */
4506                 if (mcip != NULL) {
4507                     flent = mcip->mci_flent;
4508                     ASSERT(flent->fe_rx_srs_cnt > 0);
4509
4510                 }
4511             }
4512         }
4513     }
4514
4515     /*
4516      * Set up SRS/SR according to the ring type.
4517     */
4518     switch (ring->mr_type) {
4519         case MAC_RING_TYPE_RX:
4520             /*
4521               * Setup SRS on top of the new ring if the group is
4522               * reserved for someones exclusive use.
4523             */
4524             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4525                 mac_client_impl_t *mcip;
4526
4527                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4528
4529                 /*
4530                   * Even though this group is reserved we might still
4531                   * have multiple clients, i.e a VLAN shares the
4532                   * group with the primary mac client.
4533                   */
4534                 if (mcip != NULL) {
4535                     flent = mcip->mci_flent;
4536                     ASSERT(flent->fe_rx_srs_cnt > 0);
4537
4538                 }
4539             }
4540         }
4541     }
4542
4543     /*
4544      * Set up SRS/SR according to the ring type.
4545     */
4546     switch (ring->mr_type) {
4547         case MAC_RING_TYPE_RX:
4548             /*
4549               * Setup SRS on top of the new ring if the group is
4550               * reserved for someones exclusive use.
4551             */
4552             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4553                 mac_client_impl_t *mcip;
4554
4555                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4556
4557                 /*
4558                   * Even though this group is reserved we might still
4559                   * have multiple clients, i.e a VLAN shares the
4560                   * group with the primary mac client.
4561                   */
4562                 if (mcip != NULL) {
4563                     flent = mcip->mci_flent;
4564                     ASSERT(flent->fe_rx_srs_cnt > 0);
4565
4566                 }
4567             }
4568         }
4569     }
4570
4571     /*
4572      * Set up SRS/SR according to the ring type.
4573     */
4574     switch (ring->mr_type) {
4575         case MAC_RING_TYPE_RX:
4576             /*
4577               * Setup SRS on top of the new ring if the group is
4578               * reserved for someones exclusive use.
4579             */
4580             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4581                 mac_client_impl_t *mcip;
4582
4583                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4584
4585                 /*
4586                   * Even though this group is reserved we might still
4587                   * have multiple clients, i.e a VLAN shares the
4588                   * group with the primary mac client.
4589                   */
4590                 if (mcip != NULL) {
4591                     flent = mcip->mci_flent;
4592                     ASSERT(flent->fe_rx_srs_cnt > 0);
4593
4594                 }
4595             }
4596         }
4597     }
4598
4599     /*
4600      * Set up SRS/SR according to the ring type.
4601     */
4602     switch (ring->mr_type) {
4603         case MAC_RING_TYPE_RX:
4604             /*
4605               * Setup SRS on top of the new ring if the group is
4606               * reserved for someones exclusive use.
4607             */
4608             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4609                 mac_client_impl_t *mcip;
4610
4611                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4612
4613                 /*
4614                   * Even though this group is reserved we might still
4615                   * have multiple clients, i.e a VLAN shares the
4616                   * group with the primary mac client.
4617                   */
4618                 if (mcip != NULL) {
4619                     flent = mcip->mci_flent;
4620                     ASSERT(flent->fe_rx_srs_cnt > 0);
4621
4622                 }
4623             }
4624         }
4625     }
4626
4627     /*
4628      * Set up SRS/SR according to the ring type.
4629     */
4630     switch (ring->mr_type) {
4631         case MAC_RING_TYPE_RX:
4632             /*
4633               * Setup SRS on top of the new ring if the group is
4634               * reserved for someones exclusive use.
4635             */
4636             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4637                 mac_client_impl_t *mcip;
4638
4639                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4640
4641                 /*
4642                   * Even though this group is reserved we might still
4643                   * have multiple clients, i.e a VLAN shares the
4644                   * group with the primary mac client.
4645                   */
4646                 if (mcip != NULL) {
4647                     flent = mcip->mci_flent;
4648                     ASSERT(flent->fe_rx_srs_cnt > 0);
4649
4650                 }
4651             }
4652         }
4653     }
4654
4655     /*
4656      * Set up SRS/SR according to the ring type.
4657     */
4658     switch (ring->mr_type) {
4659         case MAC_RING_TYPE_RX:
4660             /*
4661               * Setup SRS on top of the new ring if the group is
4662               * reserved for someones exclusive use.
4663             */
4664             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4665                 mac_client_impl_t *mcip;
4666
4667                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4668
4669                 /*
4670                   * Even though this group is reserved we might still
4671                   * have multiple clients, i.e a VLAN shares the
4672                   * group with the primary mac client.
4673                   */
4674                 if (mcip != NULL) {
4675                     flent = mcip->mci_flent;
4676                     ASSERT(flent->fe_rx_srs_cnt > 0);
4677
4678                 }
4679             }
4680         }
4681     }
4682
4683     /*
4684      * Set up SRS/SR according to the ring type.
4685     */
4686     switch (ring->mr_type) {
4687         case MAC_RING_TYPE_RX:
4688             /*
4689               * Setup SRS on top of the new ring if the group is
4690               * reserved for someones exclusive use.
4691             */
4692             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4693                 mac_client_impl_t *mcip;
4694
4695                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4696
4697                 /*
4698                   * Even though this group is reserved we might still
4699                   * have multiple clients, i.e a VLAN shares the
4700                   * group with the primary mac client.
4701                   */
4702                 if (mcip != NULL) {
4703                     flent = mcip->mci_flent;
4704                     ASSERT(flent->fe_rx_srs_cnt > 0);
4705
4706                 }
4707             }
4708         }
4709     }
4710
4711     /*
4712      * Set up SRS/SR according to the ring type.
4713     */
4714     switch (ring->mr_type) {
4715         case MAC_RING_TYPE_RX:
4716             /*
4717               * Setup SRS on top of the new ring if the group is
4718               * reserved for someones exclusive use.
4719             */
4720             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4721                 mac_client_impl_t *mcip;
4722
4723                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4724
4725                 /*
4726                   * Even though this group is reserved we might still
4727                   * have multiple clients, i.e a VLAN shares the
4728                   * group with the primary mac client.
4729                   */
4730                 if (mcip != NULL) {
4731                     flent = mcip->mci_flent;
4732                     ASSERT(flent->fe_rx_srs_cnt > 0);
4733
4734                 }
4735             }
4736         }
4737     }
4738
4739     /*
4740      * Set up SRS/SR according to the ring type.
4741     */
4742     switch (ring->mr_type) {
4743         case MAC_RING_TYPE_RX:
4744             /*
4745               * Setup SRS on top of the new ring if the group is
4746               * reserved for someones exclusive use.
4747             */
4748             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4749                 mac_client_impl_t *mcip;
4750
4751                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4752
4753                 /*
4754                   * Even though this group is reserved we might still
4755                   * have multiple clients, i.e a VLAN shares the
4756                   * group with the primary mac client.
4757                   */
4758                 if (mcip != NULL) {
4759                     flent = mcip->mci_flent;
4760                     ASSERT(flent->fe_rx_srs_cnt > 0);
4761
4762                 }
4763             }
4764         }
4765     }
4766
4767     /*
4768      * Set up SRS/SR according to the ring type.
4769     */
4770     switch (ring->mr_type) {
4771         case MAC_RING_TYPE_RX:
4772             /*
4773               * Setup SRS on top of the new ring if the group is
4774               * reserved for someones exclusive use.
4775             */
4776             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4777                 mac_client_impl_t *mcip;
4778
4779                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4780
4781                 /*
4782                   * Even though this group is reserved we might still
4783                   * have multiple clients, i.e a VLAN shares the
4784                   * group with the primary mac client.
4785                   */
4786                 if (mcip != NULL) {
4787                     flent = mcip->mci_flent;
4788                     ASSERT(flent->fe_rx_srs_cnt > 0);
4789
4790                 }
4791             }
4792         }
4793     }
4794
4795     /*
4796      * Set up SRS/SR according to the ring type.
4797     */
4798     switch (ring->mr_type) {
4799         case MAC_RING_TYPE_RX:
4800             /*
4801               * Setup SRS on top of the new ring if the group is
4802               * reserved for someones exclusive use.
4803             */
4804             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4805                 mac_client_impl_t *mcip;
4806
4807                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4808
4809                 /*
4810                   * Even though this group is reserved we might still
4811                   * have multiple clients, i.e a VLAN shares the
4812                   * group with the primary mac client.
4813                   */
4814                 if (mcip != NULL) {
4815                     flent = mcip->mci_flent;
4816                     ASSERT(flent->fe_rx_srs_cnt > 0);
4817
4818                 }
4819             }
4820         }
4821     }
4822
4823     /*
4824      * Set up SRS/SR according to the ring type.
4825     */
4826     switch (ring->mr_type) {
4827         case MAC_RING_TYPE_RX:
4828             /*
4829               * Setup SRS on top of the new ring if the group is
4830               * reserved for someones exclusive use.
4831             */
4832             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4833                 mac_client_impl_t *mcip;
4834
4835                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4836
4837                 /*
4838                   * Even though this group is reserved we might still
4839                   * have multiple clients, i.e a VLAN shares the
4840                   * group with the primary mac client.
4841                   */
4842                 if (mcip != NULL) {
4843                     flent = mcip->mci_flent;
4844                     ASSERT(flent->fe_rx_srs_cnt > 0);
4845
4846                 }
4847             }
4848         }
4849     }
4850
4851     /*
4852      * Set up SRS/SR according to the ring type.
4853     */
4854     switch (ring->mr_type) {
4855         case MAC_RING_TYPE_RX:
4856             /*
4857               * Setup SRS on top of the new ring if the group is
4858               * reserved for someones exclusive use.
4859             */
4860             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4861                 mac_client_impl_t *mcip;
4862
4863                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4864
4865                 /*
4866                   * Even though this group is reserved we might still
4867                   * have multiple clients, i.e a VLAN shares the
4868                   * group with the primary mac client.
4869                   */
4870                 if (mcip != NULL) {
4871                     flent = mcip->mci_flent;
4872                     ASSERT(flent->fe_rx_srs_cnt > 0);
4873
4874                 }
4875             }
4876         }
4877     }
4878
4879     /*
4880      * Set up SRS/SR according to the ring type.
4881     */
4882     switch (ring->mr_type) {
4883         case MAC_RING_TYPE_RX:
4884             /*
4885               * Setup SRS on top of the new ring if the group is
4886               * reserved for someones exclusive use.
4887             */
4888             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4889                 mac_client_impl_t *mcip;
4890
4891                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4892
4893                 /*
4894                   * Even though this group is reserved we might still
4895                   * have multiple clients, i.e a VLAN shares the
4896                   * group with the primary mac client.
4897                   */
4898                 if (mcip != NULL) {
4899                     flent = mcip->mci_flent;
4900                     ASSERT(flent->fe_rx_srs_cnt > 0);
4901
4902                 }
4903             }
4904         }
4905     }
4906
4907     /*
4908      * Set up SRS/SR according to the ring type.
4909     */
4910     switch (ring->mr_type) {
4911         case MAC_RING_TYPE_RX:
4912             /*
4913               * Setup SRS on top of the new ring if the group is
4914               * reserved for someones exclusive use.
4915             */
4916             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4917                 mac_client_impl_t *mcip;
4918
4919                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4920
4921                 /*
4922                   * Even though this group is reserved we might still
4923                   * have multiple clients, i.e a VLAN shares the
4924                   * group with the primary mac client.
4925                   */
4926                 if (mcip != NULL) {
4927                     flent = mcip->mci_flent;
4928                     ASSERT(flent->fe_rx_srs_cnt > 0);
4929
4930                 }
4931             }
4932         }
4933     }
4934
4935     /*
4936      * Set up SRS/SR according to the ring type.
4937     */
4938     switch (ring->mr_type) {
4939         case MAC_RING_TYPE_RX:
4940             /*
4941               * Setup SRS on top of the new ring if the group is
4942               * reserved for someones exclusive use.
4943             */
4944             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4945                 mac_client_impl_t *mcip;
4946
4947                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4948
4949                 /*
4950                   * Even though this group is reserved we might still
4951                   * have multiple clients, i.e a VLAN shares the
4952                   * group with the primary mac client.
4953                   */
4954                 if (mcip != NULL) {
4955                     flent = mcip->mci_flent;
4956                     ASSERT(flent->fe_rx_srs_cnt > 0);
4957
4958                 }
4959             }
4960         }
4961     }
4962
4963     /*
4964      * Set up SRS/SR according to the ring type.
4965     */
4966     switch (ring->mr_type) {
4967         case MAC_RING_TYPE_RX:
4968             /*
4969               * Setup SRS on top of the new ring if the group is
4970               * reserved for someones exclusive use.
4971             */
4972             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
4973                 mac_client_impl_t *mcip;
4974
4975                 mcip = MAC_GROUP_ONLY_CLIENT(group);
4976
4977                 /*
4978                   * Even though this group is reserved we might still
4979                   * have multiple clients, i.e a VLAN shares the
4980                   * group with the primary mac client.
4981                   */
4982                 if (mcip != NULL) {
4983                     flent = mcip->mci_flent;
4984                     ASSERT(flent->fe_rx_srs_cnt > 0);
4985
4986                 }
4987             }
4988         }
4989     }
4990
4991     /*
4992      * Set up SRS/SR according to the ring type.
4993     */
4994     switch (ring->mr_type) {
4995         case MAC_RING_TYPE_RX:
4996             /*
4997               * Setup SRS on top of the new ring if the group is
4998               * reserved for someones exclusive use.
4999             */
5000             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5001                 mac_client_impl_t *mcip;
5002
5003                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5004
5005                 /*
5006                   * Even though this group is reserved we might still
5007                   * have multiple clients, i.e a VLAN shares the
5008                   * group with the primary mac client.
5009                   */
5010                 if (mcip != NULL) {
5011                     flent = mcip->mci_flent;
5012                     ASSERT(flent->fe_rx_srs_cnt > 0);
5013
5014                 }
5015             }
5016         }
5017     }
5018
5019     /*
5020      * Set up SRS/SR according to the ring type.
5021     */
5022     switch (ring->mr_type) {
5023         case MAC_RING_TYPE_RX:
5024             /*
5025               * Setup SRS on top of the new ring if the group is
5026               * reserved for someones exclusive use.
5027             */
5028             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5029                 mac_client_impl_t *mcip;
5030
5031                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5032
5033                 /*
5034                   * Even though this group is reserved we might still
5035                   * have multiple clients, i.e a VLAN shares the
5036                   * group with the primary mac client.
5037                   */
5038                 if (mcip != NULL) {
5039                     flent = mcip->mci_flent;
5040                     ASSERT(flent->fe_rx_srs_cnt > 0);
5041
5042                 }
5043             }
5044         }
5045     }
5046
5047     /*
5048      * Set up SRS/SR according to the ring type.
5049     */
5050     switch (ring->mr_type) {
5051         case MAC_RING_TYPE_RX:
5052             /*
5053               * Setup SRS on top of the new ring if the group is
5054               * reserved for someones exclusive use.
5055             */
5056             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5057                 mac_client_impl_t *mcip;
5058
5059                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5060
5061                 /*
5062                   * Even though this group is reserved we might still
5063                   * have multiple clients, i.e a VLAN shares the
5064                   * group with the primary mac client.
5065                   */
5066                 if (mcip != NULL) {
5067                     flent = mcip->mci_flent;
5068                     ASSERT(flent->fe_rx_srs_cnt > 0);
5069
5070                 }
5071             }
5072         }
5073     }
5074
5075     /*
5076      * Set up SRS/SR according to the ring type.
5077     */
5078     switch (ring->mr_type) {
5079         case MAC_RING_TYPE_RX:
5080             /*
5081               * Setup SRS on top of the new ring if the group is
5082               * reserved for someones exclusive use.
5083             */
5084             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5085                 mac_client_impl_t *mcip;
5086
5087                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5088
5089                 /*
5090                   * Even though this group is reserved we might still
5091                   * have multiple clients, i.e a VLAN shares the
5092                   * group with the primary mac client.
5093                   */
5094                 if (mcip != NULL) {
5095                     flent = mcip->mci_flent;
5096                     ASSERT(flent->fe_rx_srs_cnt > 0);
5097
5098                 }
5099             }
5100         }
5101     }
5102
5103     /*
5104      * Set up SRS/SR according to the ring type.
5105     */
5106     switch (ring->mr_type) {
5107         case MAC_RING_TYPE_RX:
5108             /*
5109               * Setup SRS on top of the new ring if the group is
5110               * reserved for someones exclusive use.
5111             */
5112             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5113                 mac_client_impl_t *mcip;
5114
5115                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5116
5117                 /*
5118                   * Even though this group is reserved we might still
5119                   * have multiple clients, i.e a VLAN shares the
5120                   * group with the primary mac client.
5121                   */
5122                 if (mcip != NULL) {
5123                     flent = mcip->mci_flent;
5124                     ASSERT(flent->fe_rx_srs_cnt > 0);
5125
5126                 }
5127             }
5128         }
5129     }
5130
5131     /*
5132      * Set up SRS/SR according to the ring type.
5133     */
5134     switch (ring->mr_type) {
5135         case MAC_RING_TYPE_RX:
5136             /*
5137               * Setup SRS on top of the new ring if the group is
5138               * reserved for someones exclusive use.
5139             */
5140             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5141                 mac_client_impl_t *mcip;
5142
5143                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5144
5145                 /*
5146                   * Even though this group is reserved we might still
5147                   * have multiple clients, i.e a VLAN shares the
5148                   * group with the primary mac client.
5149                   */
5150                 if (mcip != NULL) {
5151                     flent = mcip->mci_flent;
5152                     ASSERT(flent->fe_rx_srs_cnt > 0);
5153
5154                 }
5155             }
5156         }
5157     }
5158
5159     /*
5160      * Set up SRS/SR according to the ring type.
5161     */
5162     switch (ring->mr_type) {
5163         case MAC_RING_TYPE_RX:
5164             /*
5165               * Setup SRS on top of the new ring if the group is
5166               * reserved for someones exclusive use.
5167             */
5168             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5169                 mac_client_impl_t *mcip;
5170
5171                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5172
5173                 /*
5174                   * Even though this group is reserved we might still
5175                   * have multiple clients, i.e a VLAN shares the
5176                   * group with the primary mac client.
5177                   */
5178                 if (mcip != NULL) {
5179                     flent = mcip->mci_flent;
5180                     ASSERT(flent->fe_rx_srs_cnt > 0);
5181
5182                 }
5183             }
5184         }
5185     }
5186
5187     /*
5188      * Set up SRS/SR according to the ring type.
5189     */
5190     switch (ring->mr_type) {
5191         case MAC_RING_TYPE_RX:
5192             /*
5193               * Setup SRS on top of the new ring if the group is
5194               * reserved for someones exclusive use.
5195             */
5196             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5197                 mac_client_impl_t *mcip;
5198
5199                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5200
5201                 /*
5202                   * Even though this group is reserved we might still
5203                   * have multiple clients, i.e a VLAN shares the
5204                   * group with the primary mac client.
5205                   */
5206                 if (mcip != NULL) {
5207                     flent = mcip->mci_flent;
5208                     ASSERT(flent->fe_rx_srs_cnt > 0);
5209
5210                 }
5211             }
5212         }
5213     }
5214
5215     /*
5216      * Set up SRS/SR according to the ring type.
5217     */
5218     switch (ring->mr_type) {
5219         case MAC_RING_TYPE_RX:
5220             /*
5221               * Setup SRS on top of the new ring if the group is
5222               * reserved for someones exclusive use.
5223             */
5224             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5225                 mac_client_impl_t *mcip;
5226
5227                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5228
5229                 /*
5230                   * Even though this group is reserved we might still
5231                   * have multiple clients, i.e a VLAN shares the
5232                   * group with the primary mac client.
5233                   */
5234                 if (mcip != NULL) {
5235                     flent = mcip->mci_flent;
5236                     ASSERT(flent->fe_rx_srs_cnt > 0);
5237
5238                 }
5239             }
5240         }
5241     }
5242
5243     /*
5244      * Set up SRS/SR according to the ring type.
5245     */
5246     switch (ring->mr_type) {
5247         case MAC_RING_TYPE_RX:
5248             /*
5249               * Setup SRS on top of the new ring if the group is
5250               * reserved for someones exclusive use.
5251             */
5252             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5253                 mac_client_impl_t *mcip;
5254
5255                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5256
5257                 /*
5258                   * Even though this group is reserved we might still
5259                   * have multiple clients, i.e a VLAN shares the
5260                   * group with the primary mac client.
5261                   */
5262                 if (mcip != NULL) {
5263                     flent = mcip->mci_flent;
5264                     ASSERT(flent->fe_rx_srs_cnt > 0);
5265
5266                 }
5267             }
5268         }
5269     }
5270
5271     /*
5272      * Set up SRS/SR according to the ring type.
5273     */
5274     switch (ring->mr_type) {
5275         case MAC_RING_TYPE_RX:
5276             /*
5277               * Setup SRS on top of the new ring if the group is
5278               * reserved for someones exclusive use.
5279             */
5280             if (group->mrg_state == MAC_GROUP_STATE_RESERVED) {
5281                 mac_client_impl_t *mcip;
5282
5283                 mcip = MAC_GROUP_ONLY_CLIENT(group);
5284
5285                 /*
5286                   * Even though this group is reserved we might still
5287                   * have multiple clients, i.e a VLAN shares the
5288                   * group with the primary mac client.
5289                   */
5290                 if (mcip != NULL) {
5291                     flent = mcip->mci_flent;
5292                     ASSERT(flent->fe_rx_srs_cnt > 0);
5293
5294                 }
5295             }
5296         }
5297     }
5298
5299     /*
5300      * Set up SRS/SR according to the ring type.
5301     */
5302     switch (ring->mr_type) {
5303         case MAC_RING_TYPE_RX:
5304             /*
5305               * Setup SRS on top of the new ring if the group is
5306               * reserved for someones exclusive use.
5307             */
5308             if (group->mrg
```

```

4423             mac_rx_srs_group_setup(mcip, flent, SRST_LINK);
4424             mac_fanout_setup(mcip, flent,
4425                             MCIP_RESOURCE_PROPS(mcip), mac_rx_deliver,
4426                             mcip, NULL, NULL);
4427         } else {
4428             ring->mr_classify_type = MAC_SW_CLASSIFIER;
4429         }
4430     }
4431     break;
4432 case MAC_RING_TYPE_TX:
4433 {
4434     mac_grp_client_t      *mgcp = group->mrg_clients;
4435     mac_client_impl_t     *mcip;
4436     mac_soft_ring_set_t   *mac_srs;
4437     mac_srs_tx_t          *tx;
4438
4439     if (MAC_GROUP_NO_CLIENT(group)) {
4440         if (ring->mr_state == MR_INUSE)
4441             mac_stop_ring(ring);
4442         ring->mr_flag = 0;
4443         break;
4444     }
4445     /*
4446     * If the rings are being moved to a group that has
4447     * clients using it, then add the new rings to the
4448     * clients SRS.
4449     */
4450     while (mgcp != NULL) {
4451         boolean_t      is_aggr;
4452
4453         mcip = mgcp->mgc_client;
4454         flent = mcip->mci_flen;
4455         is_aggr = (mcip->mci_state_flags & MCIS_IS_AGGR);
4456         mac_srs = MCIP_TX_SRS(mcip);
4457         tx = &mac_srs->srs_tx;
4458         mac_tx_client_quiesce((mac_client_handle_t)mcip);
4459         /*
4460         * If we are growing from 1 to multiple rings.
4461         */
4462         if (tx->st_mode == SRS_TX_BW ||
4463             tx->st_mode == SRS_TX_SERIALIZE ||
4464             tx->st_mode == SRS_TX_DEFAULT) {
4465             mac_ring_t      *tx_ring = tx->st_arg2;
4466
4467             tx->st_arg2 = NULL;
4468             mac_tx_srs_stat_recreate(mac_srs, B_TRUE);
4469             mac_tx_srs_add_ring(mac_srs, tx_ring);
4470             if (mac_srs->srs_type & SRST_BW_CONTROL) {
4471                 tx->st_mode = is_aggr ? SRS_TX_BW_AGGR :
4472                                         SRS_TX_BW_FANOUT;
4473             } else {
4474                 tx->st_mode = is_aggr ? SRS_TX_AGGR :
4475                                         SRS_TX_FANOUT;
4476             }
4477             tx->st_func = mac_tx_get_func(tx->st_mode);
4478         }
4479         mac_tx_srs_add_ring(mac_srs, ring);
4480         mac_fanout_setup(mcip, flent, MCIP_RESOURCE_PROPS(mcip),
4481                         mac_rx_deliver, mcip, NULL, NULL);
4482         mac_tx_client_restart((mac_client_handle_t)mcip);
4483         mgcp = mgcp->mgc_next;
4484     }
4485     break;
4486 default:
4487     ASSERT(B_FALSE);

```

```

4489     }
4490     /*
4491      * For agrgr, the default ring will be NULL to begin with. If it
4492      * is NULL, then pick the first ring that gets added as the
4493      * default ring. Any ring in an aggregation can be removed at
4494      * any time (by the user action of removing a link) and if the
4495      * current default ring gets removed, then a new one gets
4496      * picked (see i_mac_group_rem_ring()).
4497     */
4498     if (mip->mi_state_flags & MIS_IS_AGGR &&
4499         mip->mi_default_tx_ring == NULL &&
4500         ring->mr_type == MAC_RING_TYPE_TX) {
4501         mip->mi_default_tx_ring = (mac_ring_handle_t)ring;
4502     }
4503
4504     MAC_RING_UNMARK(ring, MR_INCIPIENT);
4505
4506     return (0);
4507
4508 /*
4509  * Remove a ring from it's current group. MAC internal function for dynamic
4510  * grouping.
4511  *
4512  * The caller needs to call mac_perim_enter() before calling this function.
4513  */
4514 void
4515 i_mac_group_rem_ring(mac_group_t *group, mac_ring_t *ring,
4516                       boolean_t driver_call)
4517 {
4518     mac_impl_t *mip = (mac_impl_t *)group->mrg_mh;
4519     mac_capab_rings_t *cap_rings = NULL;
4520     mac_group_type_t group_type;
4521
4522     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
4523
4524     ASSERT(mac_find_ring((mac_group_handle_t)group,
4525                           ring->mr_index) == (mac_ring_handle_t)ring);
4526     ASSERT((mac_group_t *)ring->mr_gh == group);
4527     ASSERT(ring->mr_type == group->mrg_type);
4528
4529     if (ring->mr_state == MR_INUSE)
4530         mac_stop_ring(ring);
4531     switch (ring->mr_type) {
4532     case MAC_RING_TYPE_RX:
4533         group_type = mip->mi_rx_group_type;
4534         cap_rings = &mip->mi_rx_rings_cap;
4535
4536         /*
4537          * Only hardware classified packets hold a reference to the
4538          * ring all the way up the Rx path. mac_rx_srs_remove()
4539          * will take care of quiescing the Rx path and removing the
4540          * SRS. The software classified path neither holds a reference
4541          * nor any association with the ring in mac_rx.
4542         */
4543         if (ring->mr_srs != NULL) {
4544             mac_rx_srs_remove(ring->mr_srs);
4545             ring->mr_srs = NULL;
4546         }
4547
4548         break;
4549     case MAC_RING_TYPE_TX:
4550     {
4551         mac_grp_client_t      *mgcp;
4552         mac_client_impl_t     *mcip;
4553         mac_soft_ring_set_t   *mac_srs;
4554         mac_srs_tx_t          *tx;

```

```

4555     mac_ring_t          *rem_ring;
4556     mac_group_t         *defgrp;
4557     uint_t              ring_info = 0;
4558
4559     /*
4560      * For TX this function is invoked in three
4561      * cases:
4562      *
4563      * 1) In the case of a failure during the
4564      * initial creation of a group when a share is
4565      * associated with a MAC client. So the SRS is not
4566      * yet setup, and will be setup later after the
4567      * group has been reserved and populated.
4568      *
4569      * 2) From mac_release_tx_group() when freeing
4570      * a TX SRS.
4571      *
4572      * 3) In the case of aggr, when a port gets removed,
4573      * the pseudo Tx rings that it exposed gets removed.
4574      *
4575      * In the first two cases the SRS and its soft
4576      * rings are already quiesced.
4577      */
4578     if (driver_call) {
4579         mac_client_impl_t *mcip;
4580         mac_soft_ring_set_t *mac_srs;
4581         mac_soft_ring_t *srngp;
4582         mac_srs_tx_t *srs_tx;
4583
4584         if (mip->mi_state_flags & MIS_IS_AGGR &&
4585             mip->mi_default_tx_ring ==
4586             (mac_ring_handle_t)ring) {
4587             /* pick a new default Tx ring */
4588             mip->mi_default_tx_ring =
4589                 (group->mrg_rings != ring) ?
4590                     (mac_ring_handle_t)group->mrg_rings :
4591                     (mac_ring_handle_t)(ring->mr_next);
4592
4593             /* Presently only aggr case comes here */
4594             if (group->mrg_state != MAC_GROUP_STATE_RESERVED)
4595                 break;
4596
4597             mcip = MAC_GROUP_ONLY_CLIENT(group);
4598             ASSERT(mcip != NULL);
4599             ASSERT(mcip->mc_i_state_flags & MCIS_IS_AGGR);
4600             mac_srs = MCIP_TX_SRS(mcip);
4601             ASSERT(mac_srs->srs_tx.st_mode == SRS_TX_AGGR ||
4602                   mac_srs->srs_tx.st_mode == SRS_TX_BW_AGGR);
4603             srs_tx = &mac_srs->srs_tx;
4604
4605             /*
4606              * Wakeup any callers blocked on this
4607              * Tx ring due to flow control.
4608              */
4609             srngp = srs_tx->st_soft_rings[ring->mr_index];
4610             ASSERT(srngp != NULL);
4611             mac_tx_invoke_callbacks(mcip, (mac_tx_cookie_t)srngp);
4612             mac_tx_client_quiesce((mac_client_handle_t)mcip);
4613             mac_tx_srs_del_ring(mac_srs, ring);
4614             mac_tx_client_restart((mac_client_handle_t)mcip);
4615             break;
4616
4617             ASSERT(ring != (mac_ring_t *)mip->mi_default_tx_ring);
4618             group_type = mip->mi_tx_group_type;
4619             cap_rings = &mip->mi_tx_rings_cap;
4620
4621             /*
4622              * See if we need to take it out of the MAC clients using

```

```

4621             * this group
4622             */
4623             if (MAC_GROUP_NO_CLIENT(group))
4624                 break;
4625             mgcp = group->mrg_clients;
4626             defgrp = MAC_DEFAULT_TX_GROUP(mip);
4627             while (mgcp != NULL) {
4628                 mcip = mgcp->mrgc_client;
4629                 mac_srs = MCIP_TX_SRS(mcip);
4630                 tx = &mac_srs->srs_tx;
4631                 mac_tx_client_quiesce((mac_client_handle_t)mcip);
4632
4633                 /*
4634                  * If we are here when removing rings from the
4635                  * defgroup, mac_reserve_tx_ring would have
4636                  * already deleted the ring from the MAC
4637                  * clients in the group.
4638
4639                 if (group != defgrp) {
4640                     mac_tx_invoke_callbacks(mcip,
4641                         (mac_tx_cookie_t)
4642                             mac_tx_srs_get_soft_ring(mac_srs, ring));
4643                     mac_tx_srs_del_ring(mac_srs, ring);
4644
4645                 }
4646
4647                 /*
4648                  * Additionally, if we are left with only
4649                  * one ring in the group after this, we need
4650                  * to modify the mode etc. to. (We haven't
4651                  * yet taken the ring out, so we check with 2).
4652
4653                 if (group->mrg_cur_count == 2) {
4654                     if (ring->mr_next == NULL)
4655                         rem_ring = group->mrg_rings;
4656                     else
4657                         rem_ring = ring->mr_next;
4658                     mac_tx_invoke_callbacks(mcip,
4659                         (mac_tx_cookie_t)
4660                             mac_tx_srs_get_soft_ring(mac_srs,
4661                             rem_ring));
4662                     mac_tx_srs_del_ring(mac_srs, rem_ring);
4663                     if (rem_ring->mr_state != MR_INUSE) {
4664                         (void) mac_start_ring(rem_ring);
4665
4666                     tx->st_arg2 = (void *)rem_ring;
4667                     mac_tx_srs_stat_recreate(mac_srs, B_FALSE);
4668                     ring_info = mac_hwring_getinfo(
4669                         (mac_ring_handle_t)rem_ring);
4670
4671                     /*
4672                      * We are shrinking from multiple
4673                      * to 1 ring.
4674
4675                     if (mac_srs->srs_type & SRST_BW_CONTROL) {
4676                         tx->st_mode = SRS_TX_BW;
4677                     } else if (mac_tx_serialize ||
4678                         (ring_info & MAC_RING_TX_SERIALIZE)) {
4679                         tx->st_mode = SRS_TX_SERIALIZE;
4680                     } else {
4681                         tx->st_mode = SRS_TX_DEFAULT;
4682                     }
4683                     tx->st_func = mac_tx_get_func(tx->st_mode);
4684
4685                     mac_tx_client_restart((mac_client_handle_t)mcip);
4686                     mgcp = mgcp->mrgc_next;
4687
4688                 }
4689
4690             default:

```

```

4687         ASSERT(B_FALSE);
4688     }
4689
4690     /*
4691      * Remove the ring from the group.
4692      */
4693     if (ring == group->mrg_rings)
4694         group->mrg_rings = ring->mr_next;
4695     else {
4696         mac_ring_t *pre;
4697
4698         pre = group->mrg_rings;
4699         while (pre->mr_next != ring)
4700             pre = pre->mr_next;
4701         pre->mr_next = ring->mr_next;
4702     }
4703     group->mrg_cur_count--;
4704
4705     if (!driver_call) {
4706         ASSERT(group_type == MAC_GROUP_TYPE_DYNAMIC);
4707         ASSERT(group->mrg_driver == NULL ||
4708               cap_rings->mrg_gremring != NULL);
4709
4710         /*
4711          * Remove the driver level hardware ring.
4712          */
4713         if (group->mrg_driver != NULL) {
4714             cap_rings->mrg_gremring(group->mrg_driver,
4715                                       ring->mr_driver, ring->mrg_type);
4716         }
4717     }
4718
4719     ring->mr_gh = NULL;
4720     if (driver_call)
4721         mac_ring_free(mip, ring);
4722     else
4723         ring->mr_flag = 0;
4724 }
4725
4726 /**
4727  * Move a ring to the target group. If needed, remove the ring from the group
4728  * that it currently belongs to.
4729  *
4730  * The caller need to enter MAC's perimeter by calling mac_perim_enter().
4731  */
4732 static int
4733 mac_group_mov_ring(mac_impl_t *mip, mac_group_t *d_group, mac_ring_t *ring)
4734 {
4735     mac_group_t *s_group = (mac_group_t *)ring->mr_gh;
4736     int rv;
4737
4738     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
4739     ASSERT(d_group != NULL);
4740     ASSERT(s_group->mrg_mh == d_group->mrg_mh);
4741
4742     if (s_group == d_group)
4743         return (0);
4744
4745     /*
4746      * Remove it from current group first.
4747      */
4748     if (s_group != NULL)
4749         i_mac_group_rem_ring(s_group, ring, B_FALSE);
4750
4751     /*
4752      * Add it to the new group.
4753

```

```

4753         */
4754     rv = i_mac_group_add_ring(d_group, ring, 0);
4755     if (rv != 0) {
4756         /*
4757          * Failed to add ring back to source group. If
4758          * that fails, the ring is stuck in limbo, log message.
4759          */
4760         if (i_mac_group_add_ring(s_group, ring, 0)) {
4761             cmn_err(CE_WARN, "%s: failed to move ring %p\n",
4762                     mip->mi_name, (void *)ring);
4763         }
4764     }
4765
4766     return (rv);
4767 }
4768
4769 /**
4770  * Find a MAC address according to its value.
4771  */
4772 mac_address_t *
4773 mac_find_macaddr(mac_impl_t *mip, uint8_t *mac_addr)
4774 {
4775     mac_address_t *map;
4776
4777     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
4778
4779     for (map = mip->mi_addresses; map != NULL; map = map->ma_next) {
4780         if (bcmcmp(mac_addr, map->ma_addr, map->ma_len) == 0)
4781             break;
4782     }
4783
4784     return (map);
4785 }
4786
4787 /**
4788  * Check whether the MAC address is shared by multiple clients.
4789  */
4790 boolean_t
4791 mac_check_macaddr_shared(mac_address_t *map)
4792 {
4793     ASSERT(MAC_PERIM_HELD((mac_handle_t)map->ma_mip));
4794
4795     return (map->ma_nusers > 1);
4796 }
4797
4798 /**
4799  * Remove the specified MAC address from the MAC address list and free it.
4800  */
4801 static void
4802 mac_free_macaddr(mac_address_t *map)
4803 {
4804     mac_impl_t *mip = map->ma_mip;
4805
4806     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
4807     ASSERT(mip->mi_addresses != NULL);
4808
4809     map = mac_find_macaddr(mip, map->ma_addr);
4810
4811     ASSERT(map != NULL);
4812     ASSERT(map->ma_nusers == 0);
4813
4814     if (map == mip->mi_addresses) {
4815         mip->mi_addresses = map->ma_next;
4816     } else {
4817         mac_address_t *pre;

```

```

4819         pre = mip->mi_addresses;
4820         while (pre->ma_next != map)
4821             pre = pre->ma_next;
4822         pre->ma_next = map->ma_next;
4823     }
4825     kmem_free(map, sizeof (mac_address_t));
4826 }
4828 */
4829 * Add a MAC address reference for a client. If the desired MAC address
4830 * exists, add a reference to it. Otherwise, add the new address by adding
4831 * it to a reserved group or setting promiscuous mode. Won't try different
4832 * group if the group is non-NULL, so the caller must explicitly share
4833 * default group when needed.
4834 *
4835 * Note, the primary MAC address is initialized at registration time, so
4836 * to add it to default group only need to activate it if its reference
4837 * count is still zero. Also, some drivers may not have advertised RINGS
4838 * capability.
4839 */
4840 int
4841 mac_add_macaddr(mac_impl_t *mip, mac_group_t *group, uint8_t *mac_addr,
4842     boolean_t use_hw)
4843 {
4844     mac_address_t *map;
4845     int err = 0;
4846     boolean_t allocated_map = B_FALSE;
4848     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
4850     map = mac_find_macaddr(mip, mac_addr);
4852     /*
4853      * If the new MAC address has not been added. Allocate a new one
4854      * and set it up.
4855      */
4856     if (map == NULL) {
4857         map = kmem_zalloc(sizeof (mac_address_t), KM_SLEEP);
4858         map->ma_len = mip->mi_type->mt_addr_length;
4859         bcopy(mac_addr, map->ma_addr, map->ma_len);
4860         map->ma_nusers = 0;
4861         map->ma_group = group;
4862         map->ma_mip = mip;
4864         /* add the new MAC address to the head of the address list */
4865         map->ma_next = mip->mi_addresses;
4866         mip->mi_addresses = map;
4868         allocated_map = B_TRUE;
4869     }
4871     ASSERT(map->ma_group == NULL || map->ma_group == group);
4872     if (map->ma_group == NULL)
4873         map->ma_group = group;
4875     /*
4876      * If the MAC address is already in use, simply account for the
4877      * new client.
4878      */
4879     if (map->ma_nusers++ > 0)
4880         return (0);
4882     /*
4883      * Activate this MAC address by adding it to the reserved group.
4884      */

```

```

4885     if (group != NULL) {
4886         err = mac_group_addmac(group, (const uint8_t *)mac_addr);
4887         if (err == 0) {
4888             map->ma_type = MAC_ADDRESS_TYPE_UNICAST_CLASSIFIED;
4889             return (0);
4890         }
4891     }
4893     /*
4894      * The MAC address addition failed. If the client requires a
4895      * hardware classified MAC address, fail the operation.
4896      */
4897     if (use_hw) {
4898         err = ENOSPC;
4899         goto bail;
4900     }
4902     /*
4903      * Try promiscuous mode.
4904      *
4905      * For drivers that don't advertise RINGS capability, do
4906      * nothing for the primary address.
4907      */
4908     if ((group == NULL) &&
4909         (bcmcmp(map->ma_addr, mip->mi_addr, map->ma_len) == 0)) {
4910         map->ma_type = MAC_ADDRESS_TYPE_UNICAST_CLASSIFIED;
4911         return (0);
4912     }
4914     /*
4915      * Enable promiscuous mode in order to receive traffic
4916      * to the new MAC address.
4917      */
4918     if ((err = i_mac_promisc_set(mip, B_TRUE)) == 0) {
4919         map->ma_type = MAC_ADDRESS_TYPE_UNICAST_PROMISC;
4920         return (0);
4921     }
4923     /*
4924      * Free the MAC address that could not be added. Don't free
4925      * a pre-existing address, it could have been the entry
4926      * for the primary MAC address which was pre-allocated by
4927      * mac_init_macaddr(), and which must remain on the list.
4928      */
4929     bail:
4930     map->ma_nusers--;
4931     if (allocated_map)
4932         mac_free_macaddr(map);
4933     return (err);
4934 }
4936 /*
4937  * Remove a reference to a MAC address. This may cause to remove the MAC
4938  * address from an associated group or to turn off promiscuous mode.
4939  * The caller needs to handle the failure properly.
4940 */
4941 int
4942 mac_remove_macaddr(mac_address_t *map)
4943 {
4944     mac_impl_t *mip = map->ma_mip;
4945     int err = 0;
4947     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
4949     ASSERT(map == mac_find_macaddr(mip, map->ma_addr));

```

```

4951     /*
4952      * If it's not the last client using this MAC address, only update
4953      * the MAC clients count.
4954      */
4955     if (--map->ma_nusers > 0)
4956         return (0);

4958 /*
4959  * The MAC address is no longer used by any MAC client, so remove
4960  * it from its associated group, or turn off promiscuous mode
4961  * if it was enabled for the MAC address.
4962  */
4963 switch (map->ma_type) {
4964 case MAC_ADDRESS_TYPE_UNICAST_CLASSIFIED:
4965     /*
4966      * Don't free the preset primary address for drivers that
4967      * don't advertise RINGS capability.
4968      */
4969     if (map->ma_group == NULL)
4970         return (0);

4972     err = mac_group_remmac(map->ma_group, map->ma_addr);
4973     if (err == 0)
4974         map->ma_group = NULL;
4975     break;
4976 case MAC_ADDRESS_TYPE_UNICAST_PROMISC:
4977     err = i_mac_promisc_set(mip, B_FALSE);
4978     break;
4979 default:
4980     ASSERT(B_FALSE);
4981 }
4983     if (err != 0)
4984         return (err);

4986 /*
4987  * We created MAC address for the primary one at registration, so we
4988  * won't free it here. mac_fini_macaddr() will take care of it.
4989  */
4990 if (bcmcmp(map->ma_addr, mip->mi_addr, map->ma_len) != 0)
4991     mac_free_macaddr(map);

4993     return (0);
4994 }

4996 /*
4997  * Update an existing MAC address. The caller need to make sure that the new
4998  * value has not been used.
4999 */
5000 int
5001 mac_update_macaddr(mac_address_t *map, uint8_t *mac_addr)
5002 {
5003     mac_impl_t *mip = map->ma_mip;
5004     int err = 0;

5006     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
5007     ASSERT(mac_find_macaddr(mip, mac_addr) == NULL);

5009     switch (map->ma_type) {
5010 case MAC_ADDRESS_TYPE_UNICAST_CLASSIFIED:
5011     /*
5012      * Update the primary address for drivers that are not
5013      * RINGS capable.
5014      */
5015     if (mip->mi_rx_groups == NULL) {
5016         err = mip->mi_unicst(mip->mi_driver, (const uint8_t *)

```

```

5017             mac_addr);
5018         if (err != 0)
5019             return (err);
5020         break;
5021     }

5023 /*
5024  * If this MAC address is not currently in use,
5025  * simply break out and update the value.
5026  */
5027 if (map->ma_nusers == 0)
5028     break;

5030 /*
5031  * Need to replace the MAC address associated with a group.
5032  */
5033 err = mac_group_remmac(map->ma_group, map->ma_addr);
5034 if (err != 0)
5035     return (err);

5037 err = mac_group_addmac(map->ma_group, mac_addr);

5039 /*
5040  * Failure hints hardware error. The MAC layer needs to
5041  * have error notification facility to handle this.
5042  * Now, simply try to restore the value.
5043  */
5044 if (err != 0)
5045     (void) mac_group_addmac(map->ma_group, map->ma_addr);

5047 break;
5048 case MAC_ADDRESS_TYPE_UNICAST_PROMISC:
5049     /*
5050      * Need to do nothing more if in promiscuous mode.
5051      */
5052     break;
5053 default:
5054     ASSERT(B_FALSE);
5055 }

5057 /*
5058  * Successfully replaced the MAC address.
5059  */
5060 if (err == 0)
5061     bcopy(mac_addr, map->ma_addr, map->ma_len);

5063     return (err);
5064 }

5066 /*
5067  * Freshen the MAC address with new value. Its caller must have updated the
5068  * hardware MAC address before calling this function.
5069  * This function is supposed to be used to handle the MAC address change
5070  * notification from underlying drivers.
5071 */
5072 void
5073 mac_freshen_macaddr(mac_address_t *map, uint8_t *mac_addr)
5074 {
5075     mac_impl_t *mip = map->ma_mip;
5077     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
5078     ASSERT(mac_find_macaddr(mip, mac_addr) == NULL);

5080 /*
5081  * Freshen the MAC address with new value.
5082  */

```

```

5083     bcopy(mac_addr, map->ma_addr, map->ma_len);
5084     bcopy(mac_addr, mip->mi_addr, map->ma_len);

5086     /*
5087      * Update all MAC clients that share this MAC address.
5088      */
5089     mac_unicast_update_clients(mip, map);
5090 }

5092 /*
5093  * Set up the primary MAC address.
5094  */
5095 void
5096 mac_init_macaddr(mac_impl_t *mip)
5097 {
5098     mac_address_t *map;

5100    /*
5101     * The reference count is initialized to zero, until it's really
5102     * activated.
5103     */
5104     map = kmalloc(sizeof (mac_address_t), KM_SLEEP);
5105     map->ma_len = mip->mi_type->mt_addr_length;
5106     bcopy(mip->mi_addr, map->ma_addr, map->ma_len);

5108    /*
5109     * If driver advertises RINGS capability, it shouldn't have initialized
5110     * its primary MAC address. For other drivers, including VNIC, the
5111     * primary address must work after registration.
5112     */
5113     if (mip->mi_rx_groups == NULL)
5114         map->ma_type = MAC_ADDRESS_TYPE_UNICAST_CLASSIFIED;

5116     map->ma_mip = mip;

5118     mip->mi_addresses = map;
5119 }

5121 /*
5122  * Clean up the primary MAC address. Note, only one primary MAC address
5123  * is allowed. All other MAC addresses must have been freed appropriately.
5124  */
5125 void
5126 mac_fini_macaddr(mac_impl_t *mip)
5127 {
5128     mac_address_t *map = mip->mi_addresses;

5130     if (map == NULL)
5131         return;

5133     /*
5134      * If mi_addresses is initialized, there should be exactly one
5135      * entry left on the list with no users.
5136      */
5137     ASSERT(map->ma_nusers == 0);
5138     ASSERT(map->ma_next == NULL);

5140     kmem_free(map, sizeof (mac_address_t));
5141     mip->mi_addresses = NULL;
5142 }

5144 /*
5145  * Logging related functions.
5146  *
5147  * Note that Kernel statistics have been extended to maintain fine
5148  * granularity of statistics viz. hardware lane, software lane, fanout

```

```

5149     * stats etc. However, extended accounting continues to support only
5150     * aggregate statistics like before.
5151     */

5153 /* Write the flow description to a netinfo_t record */
5154 static netinfo_t *
5155 mac_write_flow_desc(flow_entry_t *flent, mac_client_impl_t *mcip)
5156 {
5157     netinfo_t             *ninfo;
5158     net_desc_t            *ndesc;
5159     flow_desc_t           *fdesc;
5160     mac_resource_props_t *mrp;

5162     ninfo = kmalloc(sizeof (netinfo_t), KM_NOSLEEP);
5163     if (ninfo == NULL)
5164         return (NULL);
5165     ndesc = kmalloc(sizeof (net_desc_t), KM_NOSLEEP);
5166     if (ndesc == NULL) {
5167         kmem_free(ninfo, sizeof (netinfo_t));
5168         return (NULL);
5169     }

5171     /*
5172      * Grab the fe_lock to see a self-consistent fe_flow_desc.
5173      * Updates to the fe_flow_desc are done under the fe_lock
5174      */
5175     mutex_enter(&flent->fe_lock);
5176     fdesc = &flent->fe_flow_desc;
5177     mrp = &flent->fe_resource_props;

5179     ndesc->nd_name = flent->fe_flow_name;
5180     ndesc->nd_devname = mcip->mci_name;
5181     bcopy(fdesc->fd_src_mac, ndesc->nd_ehost, ETHERADDRL);
5182     bcopy(fdesc->fd_dst_mac, ndesc->nd_edest, ETHERADDRL);
5183     ndesc->nd_sap = htonl(fdesc->fd_sap);
5184     ndesc->nd_isv4 = (uint8_t)fdesc->fd_ipversion == IPV4_VERSION;
5185     ndesc->nd_bw_limit = mrp->mrp_maxbw;
5186     if (ndesc->nd_isv4) {
5187         ndesc->nd_saddr[3] = htonl(fd_local_addr.s6_addr32[3]);
5188         ndesc->nd_daddr[3] = htonl(fd_remote_addr.s6_addr32[3]);
5189     } else {
5190         bcopy(&fdesc->fd_local_addr, ndesc->nd_saddr, IPV6_ADDR_LEN);
5191         bcopy(&fdesc->fd_remote_addr, ndesc->nd_daddr, IPV6_ADDR_LEN);
5192     }
5193     ndesc->nd_sport = htons(fd_desc->fd_local_port);
5194     ndesc->nd_dport = htons(fd_desc->fd_remote_port);
5195     ndesc->nd_protocol = (uint8_t)fd_desc->fd_protocol;
5196     mutex_exit(&flent->fe_lock);

5198     ninfo->ni_record = ndesc;
5199     ninfo->ni_size = sizeof (net_desc_t);
5200     ninfo->ni_type = EX_NET_FLDESC_REC;
5202 }
5203 }

5205 /* Write the flow statistics to a netinfo_t record */
5206 static netinfo_t *
5207 mac_write_flow_stats(flow_entry_t *flent)
5208 {
5209     netinfo_t             *ninfo;
5210     net_stat_t            *nstat;
5211     mac_soft_ring_set_t  *mac_srs;
5212     mac_rx_stats_t        *mac_rx_stat;
5213     mac_tx_stats_t        *mac_tx_stat;
5214     int                   i;

```

```

5216     ninfo = kmem_zalloc(sizeof (netinfo_t), KM_NOSLEEP);
5217     if (ninfo == NULL)
5218         return (NULL);
5219     nstat = kmem_zalloc(sizeof (net_stat_t), KM_NOSLEEP);
5220     if (nstat == NULL) {
5221         kmem_free(ninfo, sizeof (netinfo_t));
5222         return (NULL);
5223     }
5224
5225     nstat->ns_name = flent->fe_flow_name;
5226     for (i = 0; i < flent->fe_rx_srs_cnt; i++) {
5227         mac_srs = (mac_soft_ring_set_t *)flent->fe_rx_srs[i];
5228         mac_rx_stat = &mac_srs->srs_rx.sr_stat;
5229
5230         nstat->ns_ibytes += mac_rx_stat->mrs_intrbytes +
5231             mac_rx_stat->mrs_pollbytes + mac_rx_stat->mrs_lclbytes;
5232         nstat->ns_ipackets += mac_rx_stat->mrs_intrcnt +
5233             mac_rx_stat->mrs_pollcnt + mac_rx_stat->mrs_lclcnt;
5234         nstat->ns_oerrors += mac_rx_stat->mrs_ierrors;
5235     }
5236
5237     mac_srs = (mac_soft_ring_set_t *)(flent->fe_tx_srs);
5238     if (mac_srs != NULL) {
5239         mac_tx_stat = &mac_srs->srs_tx.st_stat;
5240
5241         nstat->ns_obytes = mac_tx_stat->mts_obytes;
5242         nstat->ns_opackets = mac_tx_stat->mts_opackets;
5243         nstat->ns_oerrors = mac_tx_stat->mts_oerrors;
5244     }
5245
5246     ninfo->ni_record = nstat;
5247     ninfo->ni_size = sizeof (net_stat_t);
5248     ninfo->ni_type = EX_NET_FLSTAT_REC;
5249
5250     return (ninfo);
5251 }
5252
5253 /* Write the link description to a netinfo_t record */
5254 static netinfo_t *
5255 mac_write_link_desc(mac_client_impl_t *mcip)
5256 {
5257     netinfo_t          *ninfo;
5258     net_desc_t          *ndesc;
5259     flow_entry_t        *flent = mcip->mci_flent;
5260
5261     ninfo = kmem_zalloc(sizeof (netinfo_t), KM_NOSLEEP);
5262     if (ninfo == NULL)
5263         return (NULL);
5264     ndesc = kmem_zalloc(sizeof (net_desc_t), KM_NOSLEEP);
5265     if (ndesc == NULL) {
5266         kmem_free(ninfo, sizeof (netinfo_t));
5267         return (NULL);
5268     }
5269
5270     ndesc->nd_name = mcip->mci_name;
5271     ndesc->nd_devname = mcip->mci_name;
5272     ndesc->nd_isv4 = B_TRUE;
5273
5274     /* Grab the fe_lock to see a self-consistent fe_flow_desc.
5275      * Updates to the fe_flow_desc are done under the fe_lock
5276      * after removing the flent from the flow table.
5277     */
5278     mutex_enter(&flent->fe_lock);
5279     bcopy(fluent->fe_flow_desc.fd_src_mac, ndesc->nd_ehost, ETHERADDRLEN);
5280     mutex_exit(&flent->fe_lock);

```

```

5282     ninfo->ni_record = ndesc;
5283     ninfo->ni_size = sizeof (net_desc_t);
5284     ninfo->ni_type = EX_NET_LNDESC_REC;
5285
5286     return (ninfo);
5287 }
5288
5289 /* Write the link statistics to a netinfo_t record */
5290 static netinfo_t *
5291 mac_write_link_stats(mac_client_impl_t *mcip)
5292 {
5293     netinfo_t          *ninfo;
5294     net_stat_t          *nstat;
5295     flow_entry_t        *flent;
5296     mac_soft_ring_set_t *mac_srs;
5297     mac_rx_stats_t      *mac_rx_stat;
5298     mac_tx_stats_t      *mac_tx_stat;
5299     int                  i;
5300
5301     ninfo = kmem_zalloc(sizeof (netinfo_t), KM_NOSLEEP);
5302     if (ninfo == NULL)
5303         return (NULL);
5304     nstat = kmem_zalloc(sizeof (net_stat_t), KM_NOSLEEP);
5305     if (nstat == NULL) {
5306         kmem_free(ninfo, sizeof (netinfo_t));
5307         return (NULL);
5308     }
5309
5310     nstat->ns_name = mcip->mci_name;
5311     flent = mcip->mci_flent;
5312     if (flent != NULL) {
5313         for (i = 0; i < flent->fe_rx_srs_cnt; i++) {
5314             mac_srs = (mac_soft_ring_set_t *)flent->fe_rx_srs[i];
5315             mac_rx_stat = &mac_srs->srs_rx.sr_stat;
5316
5317             nstat->ns_ibytes += mac_rx_stat->mrs_intrbytes +
5318                 mac_rx_stat->mrs_pollbytes +
5319                 mac_rx_stat->mrs_lclbytes;
5320             nstat->ns_ipackets += mac_rx_stat->mrs_intrcnt +
5321                 mac_rx_stat->mrs_pollcnt + mac_rx_stat->mrs_lclcnt;
5322             nstat->ns_oerrors += mac_rx_stat->mrs_ierrors;
5323         }
5324     }
5325
5326     mac_srs = (mac_soft_ring_set_t *) (mcip->mci_flent->fe_tx_srs);
5327     if (mac_srs != NULL) {
5328         mac_tx_stat = &mac_srs->srs_tx.st_stat;
5329
5330         nstat->ns_obytes = mac_tx_stat->mts_obytes;
5331         nstat->ns_opackets = mac_tx_stat->mts_opackets;
5332         nstat->ns_oerrors = mac_tx_stat->mts_oerrors;
5333     }
5334
5335     ninfo->ni_record = nstat;
5336     ninfo->ni_size = sizeof (net_stat_t);
5337     ninfo->ni_type = EX_NET_LNSTAT_REC;
5338
5339     return (ninfo);
5340 }
5341
5342 typedef struct i_mac_log_state_s {
5343     boolean_t          mi_last;
5344     int                mi_fenable;
5345     int                mi_lenable;
5346     list_t             *mi_list;

```

```

5347 } i_mac_log_state_t;

5349 /*
5350 * For a given flow, if the description has not been logged before, do it now.
5351 * If it is a VNIC, then we have collected information about it from the MAC
5352 * table, so skip it.
5353 *
5354 * Called through mac_flow_walk_nolock()
5355 *
5356 * Return 0 if successful.
5357 */
5358 static int
5359 mac_log_flowinfo(flow_entry_t *flemt, void *arg)
5360 {
5361     mac_client_impl_t      *mcip = flemt->fe_mcip;
5362     i_mac_log_state_t      *lstate = arg;
5363     netinfo_t               *ninfo;
5364
5365     if (mcip == NULL)
5366         return (0);
5367
5368     /*
5369     * If the name starts with "vnic", and fe_user_generated is true (to
5370     * exclude the mcast and active flow entries created implicitly for
5371     * a vnic, it is a VNIC flow. i.e. vnicl is a vnic flow,
5372     * vnic/bgel/mcastl is not and neither is vnic/bgel/active.
5373     */
5374     if (strncasecmp(flemt->fe_flow_name, "vnic", 4) == 0 &&
5375         (flemt->fe_type & FLOW_USER) != 0) {
5376         return (0);
5377     }
5378
5379     if (!flemt->fe_desc_logged) {
5380         /*
5381         * We don't return error because we want to continue the
5382         * walk in case this is the last walk which means we
5383         * need to reset fe_desc_logged in all the flows.
5384         */
5385         if ((ninfo = mac_write_flow_desc(flemt, mcip)) == NULL)
5386             return (0);
5387         list_insert_tail(lstate->mi_list, ninfo);
5388         flemt->fe_desc_logged = B_TRUE;
5389     }
5390
5391     /*
5392     * Regardless of the error, we want to proceed in case we have to
5393     * reset fe_desc_logged.
5394     */
5395     ninfo = mac_write_flow_stats(flemt);
5396     if (ninfo == NULL)
5397         return (-1);
5398
5399     list_insert_tail(lstate->mi_list, ninfo);
5400
5401     if (mcip != NULL && !(mcip->mci_state_flags & MCIS_DESC_LOGGED))
5402         flemt->fe_desc_logged = B_FALSE;
5403
5404     return (0);
5405 }

5406 /*
5407 * Log the description for each mac client of this mac_impl_t, if it
5408 * hasn't already been done. Additionally, log statistics for the link as
5409 * well. Walk the flow table and log information for each flow as well.
5410 * If it is the last walk (mci_last), then we turn off mci_desc_logged (and
5411 * also fe_desc_logged, if flow logging is on) since we want to log the

```

```

5413     * description if and when logging is restarted.
5414     *
5415     * Return 0 upon success or -1 upon failure
5416     */
5417     static int
5418     i_mac_impl_log(mac_impl_t *mip, i_mac_log_state_t *lstate)
5419     {
5420         mac_client_impl_t      *mcip;
5421         netinfo_t               *ninfo;
5422
5423         i_mac_perim_enter(mip);
5424         /*
5425         * Only walk the client list for NIC and etherstub
5426         */
5427         if ((mip->mi_state_flags & MIS_DISABLED) ||
5428             ((mip->mi_state_flags & MIS_IS_VNIC) &&
5429             (mac_get_lower_mac_handle((mac_handle_t)mip) != NULL))) {
5430             i_mac_perim_exit(mip);
5431             return (0);
5432         }
5433
5434         for (mcip = mip->mi_clients_list; mcip != NULL;
5435             mcip = mcip->mci_client_next) {
5436             if (!MCIP_DATAPATH_SETUP(mcip))
5437                 continue;
5438             if (lstate->mi_lenable) {
5439                 if (!(mcip->mci_state_flags & MCIS_DESC_LOGGED)) {
5440                     ninfo = mac_write_link_desc(mcip);
5441                     if (ninfo == NULL) {
5442                         /*
5443                         * We can't terminate it if this is the last
5444                         * walk, else there might be some links with
5445                         * mi_desc_logged set to true, which means
5446                         * their description won't be logged the next
5447                         * time logging is started (similarly for the
5448                         * flows within such links). We can continue
5449                         * without walking the flow table (i.e. to
5450                         * set fe_desc_logged to false) because we
5451                         * won't have written any flow stuff for this
5452                         * link as we haven't logged the link itself.
5453                         */
5454                     i_mac_perim_exit(mip);
5455                     if (lstate->mi_last)
5456                         return (0);
5457                 }
5458             }
5459             mcip->mci_state_flags |= MCIS_DESC_LOGGED;
5460             list_insert_tail(lstate->mi_list, ninfo);
5461         }
5462
5463         ninfo = mac_write_link_stats(mcip);
5464         if (ninfo == NULL && !lstate->mi_last) {
5465             i_mac_perim_exit(mip);
5466             return (-1);
5467         }
5468         list_insert_tail(lstate->mi_list, ninfo);
5469
5470         if (lstate->mi_last)
5471             mcip->mci_state_flags &= ~MCIS_DESC_LOGGED;
5472
5473         if (lstate->mi_fenable) {
5474             if (mcip->mci_subflow_tab != NULL) {
5475                 (void) mac_flow_walk_nolock(
5476                     mcip->mci_subflow_tab, mac_log_flowinfo,
5477

```

```

5479
5480         }
5481     }
5482     i_mac_perim_exit(mip);
5483     return (0);
5485 }

5487 /*
5488  * modhash walker function to add a mac_impl_t to a list
5489  */
5490 /*ARGSUSED*/
5491 static uint_t
5492 i_mac_impl_list_walker(mod_hash_key_t key, mod_hash_val_t *val, void *arg)
5493 {
5494     list_t          *list = (list_t *)arg;
5495     mac_impl_t      *mip = (mac_impl_t *)val;
5496
5497     if ((mip->mi_state_flags & MIS_DISABLED) == 0) {
5498         list_insert_tail(list, mip);
5499         mip->mi_ref++;
5500     }
5502
5503     return (MH_WALK_CONTINUE);
5504 }

5505 void
5506 i_mac_log_info(list_t *net_log_list, i_mac_log_state_t *lstate)
5507 {
5508     list_t          mac_impl_list;
5509     mac_impl_t      *mip;
5510     netinfo_t        *ninfo;
5511
5512     /* Create list of mac_impls */
5513     ASSERT(RW_LOCK_HELD(&i_mac_impl_lock));
5514     list_create(&mac_impl_list, sizeof(mac_impl_t), offsetof(mac_impl_t,
5515             mi_node));
5516     mod_hash_walk(i_mac_impl_hash, i_mac_impl_list_walker, &mac_impl_list);
5517     rw_exit(&i_mac_impl_lock);
5518
5519     /* Create log entries for each mac_impl */
5520     for (mip = list_head(&mac_impl_list); mip != NULL;
5521         mip = list_next(&mac_impl_list, mip)) {
5522         if (i_mac_impl_log(mip, lstate) != 0)
5523             continue;
5524     }
5525
5526     /* Remove elements and destroy list of mac_impls */
5527     rw_enter(&i_mac_impl_lock, RW_WRITER);
5528     while ((mip = list_remove_tail(&mac_impl_list)) != NULL) {
5529         mip->mi_ref--;
5530     }
5531     rw_exit(&i_mac_impl_lock);
5532     list_destroy(&mac_impl_list);
5533
5534     /*
5535      * Write log entries to files outside of locks, free associated
5536      * structures, and remove entries from the list.
5537      */
5538     while ((ninfo = list_head(net_log_list)) != NULL) {
5539         (void) exact_commit_netinfo(ninfo->ni_record, ninfo->ni_type);
5540         list_remove(net_log_list, ninfo);
5541         kmem_free(ninfo->ni_record, ninfo->ni_size);
5542         kmem_free(ninfo, sizeof(*ninfo));
5543     }
5544     list_destroy(net_log_list);

```

```

5545 }

5547 /*
5548  * The timer thread that runs every mac_logging_interval seconds and logs
5549  * link and/or flow information.
5550  */
5551 /* ARGSUSED */
5552 void
5553 mac_log_linkinfo(void *arg)
5554 {
5555     i_mac_log_state_t      lstate;
5556     list_t                 net_log_list;
5557
5558     list_create(&net_log_list, sizeof(netinfo_t),
5559                 offsetof(netinfo_t, ni_link));
5560
5561     rw_enter(&i_mac_impl_lock, RW_READER);
5562     if (!mac_flow_log_enable && !mac_link_log_enable) {
5563         rw_exit(&i_mac_impl_lock);
5564         return;
5565     }
5566     lstate.mi_fenable = mac_flow_log_enable;
5567     lstate.mi_lenable = mac_link_log_enable;
5568     lstate.mi_last = B_FALSE;
5569     lstate.mi_list = &net_log_list;
5570
5571     /* Write log entries for each mac_impl in the list */
5572     i_mac_log_info(&net_log_list, &lstate);
5573
5574     if (mac_flow_log_enable || mac_link_log_enable) {
5575         mac_logging_timer = timeout(mac_log_linkinfo, NULL,
5576                                     SEC_TO_TICK(mac_logging_interval));
5577     }
5578 }

5579 typedef struct i_mac_fastpath_state_s {
5580     boolean_t      mf_disable;
5581     int           mf_err;
5582 } i_mac_fastpath_state_t;

5583 /* modhash walker function to enable or disable fastpath */
5584 /*ARGSUSED*/
5585 static uint_t
5586 i_mac_fastpath_walker(mod_hash_key_t key, mod_hash_val_t *val,
5587                         void *arg)
5588 {
5589     i_mac_fastpath_state_t *state = arg;
5590     mac_handle_t            mh = (mac_handle_t)val;
5591
5592     if (state->mf_disable)
5593         state->mf_err = mac_fastpath_disable(mh);
5594     else
5595         mac_fastpath_enable(mh);
5596
5597     return (state->mf_err == 0 ? MH_WALK_CONTINUE : MH_WALK_TERMINATE);
5598 }

5599 /*
5600  * Start the logging timer.
5601  */
5602 int
5603 mac_start_logusage(mac_logtype_t type, uint_t interval)
5604 {
5605     i_mac_fastpath_state_t dstate = {B_TRUE, 0};
5606     i_mac_fastpath_state_t estate = {B_FALSE, 0};
5607
5608     i_mac_fastpath_state_t err;
5609
5610     int

```

```

5612     rw_enter(&i_mac_impl_lock, RW_WRITER);
5613     switch (type) {
5614         case MAC_LOGTYPE_FLOW:
5615             if (mac_flow_log_enable) {
5616                 rw_exit(&i_mac_impl_lock);
5617                 return (0);
5618             }
5619             /* FALLTHRU */
5620         case MAC_LOGTYPE_LINK:
5621             if (mac_link_log_enable) {
5622                 rw_exit(&i_mac_impl_lock);
5623                 return (0);
5624             }
5625             break;
5626         default:
5627             ASSERT(0);
5628     }
5629
5630     /* Disable fastpath */
5631     mod_hash_walk(i_mac_impl_hash, i_mac_fastpath_walker, &dstate);
5632     if ((err = dstate.mf_err) != 0) {
5633         /* Reenable fastpath */
5634         mod_hash_walk(i_mac_impl_hash, i_mac_fastpath_walker, &estate);
5635         rw_exit(&i_mac_impl_lock);
5636         return (err);
5637     }
5638
5639     switch (type) {
5640         case MAC_LOGTYPE_FLOW:
5641             mac_flow_log_enable = B_TRUE;
5642             /* FALLTHRU */
5643         case MAC_LOGTYPE_LINK:
5644             mac_link_log_enable = B_TRUE;
5645             break;
5646     }
5647
5648     mac_logging_interval = interval;
5649     rw_exit(&i_mac_impl_lock);
5650     mac_log_linkinfo(NULL);
5651     return (0);
5652 }
5653 */
5654 * Stop the logging timer if both link and flow logging are turned off.
5655 */
5656 void
5657 mac_stop_logusage(mac_logtype_t type)
5658 {
5659     i_mac_log_state_t lstate;
5660     i_mac_fastpath_state_t estate = {B_FALSE, 0};
5661     list_t net_log_list;
5662
5663     list_create(&net_log_list, sizeof (netinfo_t),
5664                 offsetof(netinfo_t, ni_link));
5665
5666     rw_enter(&i_mac_impl_lock, RW_WRITER);
5667
5668     lstate.mi_fenable = mac_flow_log_enable;
5669     lstate.mi_lenable = mac_link_log_enable;
5670     lstate.mi_list = &net_log_list;
5671
5672     /* Last walk */
5673     lstate.mi_last = B_TRUE;
5674
5675     switch (type) {

```

```

5676         case MAC_LOGTYPE_FLOW:
5677             if (lstate.mi_fenable) {
5678                 ASSERT(mac_link_log_enable);
5679                 mac_flow_log_enable = B_FALSE;
5680                 mac_link_log_enable = B_FALSE;
5681                 break;
5682             }
5683             /* FALLTHRU */
5684         case MAC_LOGTYPE_LINK:
5685             if (!lstate.mi_lenable || mac_flow_log_enable) {
5686                 rw_exit(&i_mac_impl_lock);
5687                 return;
5688             }
5689             mac_link_log_enable = B_FALSE;
5690             break;
5691         default:
5692             ASSERT(0);
5693     }
5694
5695     /* Reenable fastpath */
5696     mod_hash_walk(i_mac_impl_hash, i_mac_fastpath_walker, &estate);
5697
5698     (void) untimout(mac_logging_timer);
5699     mac_logging_timer = 0;
5700
5701     /* Write log entries for each mac_impl in the list */
5702     i_mac_log_info(&net_log_list, &lstate);
5703
5704 }
5705 /*
5706  * Walk the rx and tx SRS/SRs for a flow and update the priority value.
5707  */
5708 void
5709 mac_flow_update_priority(mac_client_impl_t *mcip, flow_entry_t *flext)
5710 {
5711     pri_t pri;
5712     int count;
5713     mac_soft_ring_set_t *mac_srs;
5714
5715     if (flext->fe_rx_srs_cnt <= 0)
5716         return;
5717
5718     if (((mac_soft_ring_set_t *)flext->fe_rx_srs[0])->srs_type ==
5719         SRST_FLOW) {
5720         pri = FLOW_PRIORITY(mcip->mci_min_pri,
5721                             mcip->mci_max_pri,
5722                             flext->fe_resource_props.mrp_priority);
5723     } else {
5724         pri = mcip->mci_max_pri;
5725     }
5726
5727     for (count = 0; count < flext->fe_rx_srs_cnt; count++) {
5728         mac_srs = flext->fe_rx_srs[count];
5729         mac_update_srs_priority(mac_srs, pri);
5730     }
5731
5732     /*
5733      * If we have a Tx SRS, we need to modify all the threads associated
5734      * with it.
5735      */
5736     if (flext->fe_tx_srs != NULL)
5737         mac_update_srs_priority(flext->fe_tx_srs, pri);
5738 }
5739 /*
5740  * RX and TX rings are reserved according to different semantics depending
5741  * on the requests from the MAC clients and type of rings.

```

```

5743 /*
5744 * On the Tx side, by default we reserve individual rings, independently from
5745 * the groups.
5746 */
5747 /* On the Rx side, the reservation is at the granularity of the group
5748 * of rings, and used for v12n level 1 only. It has a special case for the
5749 * primary client.
5750 */
5751 /* If a share is allocated to a MAC client, we allocate a TX group and an
5752 * RX group to the client, and assign TX rings and RX rings to these
5753 * groups according to information gathered from the driver through
5754 * the share capability.
5755 */
5756 /* The foreseeable evolution of Rx rings will handle v12n level 2 and higher
5757 * to allocate individual rings out of a group and program the hw classifier
5758 * based on IP address or higher level criteria.
5759 */
5760 /*
5761 * mac_reserve_tx_ring()
5762 * Reserve a unused ring by marking it with MR_INUSE state.
5763 * As reserved, the ring is ready to function.
5764 */
5765 /*
5766 * Notes for Hybrid I/O:
5767 */
5768 /* If a specific ring is needed, it is specified through the desired_ring
5769 * argument. Otherwise that argument is set to NULL.
5770 * If the desired ring was previously allocated to another client, this
5771 * function swaps it with a new ring from the group of unassigned rings.
5772 */
5773 mac_ring_t *
5774 mac_reserve_tx_ring(mac_impl_t *mip, mac_ring_t *desired_ring)
5775 {
5776     mac_group_t           *group;
5777     mac_grp_client_t      *mgcp;
5778     mac_client_impl_t      *mcip;
5779     mac_soft_ring_set_t    *srs;
5780
5781     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
5782
5783     /*
5784     * Find an available ring and start it before changing its status.
5785     * The unassigned rings are at the end of the mi_tx_groups
5786     * array.
5787     */
5788     group = MAC_DEFAULT_TX_GROUP(mip);
5789
5790     /* Can't take the default ring out of the default group */
5791     ASSERT(desired_ring != (mac_ring_t *)mip->mi_default_tx_ring);
5792
5793     if (desired_ring->mr_state == MR_FREE) {
5794         ASSERT(MAC_GROUP_NO_CLIENT(group));
5795         if (mac_start_ring(desired_ring) != 0)
5796             return (NULL);
5797         return (desired_ring);
5798     }
5799
5800     /*
5801     * There are clients using this ring, so let's move the clients
5802     * away from using this ring.
5803     */
5804     for (mgcp = group->mrg_clients; mgcp != NULL; mgcp = mgcp->mrg_next) {
5805         mcip = mgcp->mrg_client;
5806         mac_tx_client_quiesce((mac_client_handle_t)mcip);
5807         srs = MCIP_TX_SRS(mcip);
5808         ASSERT(mac_tx_srs_ring_present(srs, desired_ring));
5809         mac_tx_invoke_callbacks(mcip,

```

```

5810             (mac_tx_cookie_t)mac_tx_srs_get_soft_ring(srs,
5811                                         desired_ring));
5812             mac_tx_srs_del_ring(srs, desired_ring);
5813             mac_tx_client_restart((mac_client_handle_t)mcip);
5814         }
5815     }
5816
5817     /*
5818     * For a reserved group with multiple clients, return the primary client.
5819     */
5820     static mac_client_impl_t *
5821     mac_get_grp_primary(mac_group_t *grp)
5822     {
5823         mac_grp_client_t          *mgcp = grp->mrg_clients;
5824         mac_client_impl_t          *mcip;
5825
5826         while (mgcp != NULL) {
5827             mcip = mgcp->mrg_client;
5828             if (mcip->mci_flext->fe_type & FLOW_PRIMARY_MAC)
5829                 return (mcip);
5830             mgcp = mgcp->mrg_next;
5831         }
5832     }
5833
5834
5835     /*
5836     * Hybrid I/O specifies the ring that should be given to a share.
5837     * If the ring is already used by clients, then we need to release
5838     * the ring back to the default group so that we can give it to
5839     * the share. This means the clients using this ring now get a
5840     * replacement ring. If there aren't any replacement rings, this
5841     * function returns a failure.
5842     */
5843     static int
5844     mac_reclaim_ring_from_grp(mac_impl_t *mip, mac_ring_type_t ring_type,
5845                               mac_ring_t *ring, mac_ring_t **rings, int nrings)
5846     {
5847         mac_group_t           *group = (mac_group_t *)ring->mr_gh;
5848         mac_resource_props_t  *mrp;
5849         mac_client_impl_t      *mcip;
5850         mac_group_t           *defgrp;
5851         mac_ring_t            *tring;
5852         mac_group_t           *tgrp;
5853         int                   i;
5854         int                   j;
5855
5856         mcip = MAC_GROUP_ONLY_CLIENT(group);
5857         if (mcip == NULL)
5858             mcip = mac_get_grp_primary(group);
5859         ASSERT(mcip != NULL);
5860         ASSERT(mcip->mci_share == NULL);
5861
5862         mrp = MCIP_RESOURCE_PROPS(mcip);
5863         if (ring_type == MAC_RING_TYPE_RX) {
5864             defgrp = mip->mi_rx_donor_grp;
5865             if ((mrp->mrp_mask & MRP_RX_RINGS) == 0) {
5866                 /* Need to put this mac client in the default group */
5867                 if (mac_rx_switch_group(mcip, group, defgrp) != 0)
5868                     return (ENOSPC);
5869             } else {
5870                 /*
5871                 * Switch this ring with some other ring from
5872                 * the default group.
5873                 */
5874                 for (tring = defgrp->mrg_rings; tring != NULL;

```

```

5875         tring = tring->mr_next) {
5876             if (tring->mr_index == 0)
5877                 continue;
5878             for (j = 0; j < nrings; j++) {
5879                 if (rings[j] == tring)
5880                     break;
5881             }
5882             if (j >= nrings)
5883                 break;
5884         }
5885         if (tring == NULL)
5886             return (ENOSPC);
5887         if (mac_group_mov_ring(mip, group, tring) != 0)
5888             return (ENOSPC);
5889         if (mac_group_mov_ring(mip, defgrp, ring) != 0) {
5890             (void) mac_group_mov_ring(mip, defgrp, tring);
5891             return (ENOSPC);
5892         }
5893     }
5894     ASSERT(ring->mr_gh == (mac_group_handle_t)defgrp);
5895     return (0);
5896 }

5897 defgrp = MAC_DEFAULT_TX_GROUP(mip);
5898 if (ring == (mac_ring_t *)mip->mi_default_tx_ring) {
5899     /*
5900      * See if we can get a spare ring to replace the default
5901      * ring.
5902     */
5903     if (defgrp->mrg_cur_count == 1) {
5904         /*
5905          * Need to get a ring from another client, see if
5906          * there are any clients that can be moved to
5907          * the default group, thereby freeing some rings.
5908         */
5909         for (i = 0; i < mip->mi_tx_group_count; i++) {
5910             tgrp = &mip->mi_tx_groups[i];
5911             if (tgrp->mrg_state ==
5912                 MAC_GROUP_STATE_REGISTERED) {
5913                 continue;
5914             }
5915             mcip = MAC_GROUP_ONLY_CLIENT(tgrp);
5916             if (mcip == NULL)
5917                 mcip = mac_get_grp_primary(tgrp);
5918             ASSERT(mcip != NULL);
5919             mrp = MCIP_RESOURCE_PROPS(mcip);
5920             if ((mrp->mrp_mask & MRP_TX_RINGS) == 0) {
5921                 ASSERT(tgrp->mrg_cur_count == 1);
5922                 /*
5923                  * If this ring is part of the
5924                  * rings asked by the share we cannot
5925                  * use it as the default ring.
5926                 */
5927                 for (j = 0; j < nrings; j++) {
5928                     if (rings[j] == tgrp->mrg_rings)
5929                         break;
5930                 }
5931                 if (j < nrings)
5932                     continue;
5933                 mac_tx_client_quiesce(
5934                     (mac_client_handle_t)mcip);
5935                 mac_tx_switch_group(mcip, tgrp,
5936                     defgrp);
5937                 mac_tx_client_restart(
5938                     (mac_client_handle_t)mcip);
5939             }
5940         }

```

```

5941         }
5942     }
5943     /*
5944      * All the rings are reserved, can't give up the
5945      * default ring.
5946     */
5947     if (defgrp->mrg_cur_count <= 1)
5948         return (ENOSPC);
5949 }
5950 /*
5951  * Swap the default ring with another.
5952 */
5953 for (tring = defgrp->mrg_rings; tring != NULL;
5954     tring = tring->mr_next) {
5955     /*
5956      * If this ring is part of the rings asked by the
5957      * share we cannot use it as the default ring.
5958     */
5959     for (j = 0; j < nrings; j++) {
5960         if (rings[j] == tring)
5961             break;
5962     }
5963     if (j >= nrings)
5964         break;
5965 }
5966 ASSERT(tring != NULL);
5967 mip->mi_default_tx_ring = (mac_ring_handle_t)tring;
5968 return (0);
5969 }
5970 /*
5971  * The Tx ring is with a group reserved by a MAC client. See if
5972  * we can swap it.
5973 */
5974 ASSERT(group->mrg_state == MAC_GROUP_STATE_RESERVED);
5975 mcip = MAC_GROUP_ONLY_CLIENT(group);
5976 if (mcip == NULL)
5977     mcip = mac_get_grp_primary(group);
5978 ASSERT(mcip != NULL);
5979 mrp = MCIP_RESOURCE_PROPS(mcip);
5980 mac_tx_client_quiesce((mac_client_handle_t)mcip);
5981 if ((mrp->mrp_mask & MRP_TX_RINGS) == 0) {
5982     ASSERT(group->mrg_cur_count == 1);
5983     /* Put this mac client in the default group */
5984     mac_tx_switch_group(mcip, group, defgrp);
5985 } else {
5986     /*
5987      * Switch this ring with some other ring from
5988      * the default group.
5989     */
5990     for (tring = defgrp->mrg_rings; tring != NULL;
5991         tring = tring->mr_next) {
5992         if (tring == (mac_ring_t *)mip->mi_default_tx_ring)
5993             continue;
5994         /*
5995          * If this ring is part of the rings asked by the
5996          * share we cannot use it for swapping.
5997         */
5998         for (j = 0; j < nrings; j++) {
5999             if (rings[j] == tring)
6000                 break;
6001         }
6002         if (j >= nrings)
6003             break;
6004     }
6005     if (tring == NULL) {
6006         mac_tx_client_restart((mac_client_handle_t)mcip);
6007     }

```

```

6007     return (ENOSPC);
6008 }
6009     if (mac_group_mov_ring(mip, group, tring) != 0) {
6010         mac_tx_client_restart((mac_client_handle_t)mcip);
6011         return (ENOSPC);
6012     }
6013     if (mac_group_mov_ring(mip, defgrp, ring) != 0) {
6014         (void) mac_group_mov_ring(mip, defgrp, tring);
6015         mac_tx_client_restart((mac_client_handle_t)mcip);
6016         return (ENOSPC);
6017     }
6018 }
6019     mac_tx_client_restart((mac_client_handle_t)mcip);
6020     ASSERT(ring->mr_gh == (mac_group_handle_t)defgrp);
6021     return (0);
6022 }

6024 /*
6025 * Populate a zero-ring group with rings. If the share is non-NULL,
6026 * the rings are chosen according to that share.
6027 * Invoked after allocating a new RX or TX group through
6028 * mac_reserve_rx_group() or mac_reserve_tx_group(), respectively.
6029 * Returns zero on success, an errno otherwise.
6030 */
6031 int
6032 i_mac_group_allocate_rings(mac_impl_t *mip, mac_ring_type_t ring_type,
6033     mac_group_t *src_group, mac_group_t *new_group, mac_share_handle_t share,
6034     uint32_t ringcnt)
6035 {
6036     mac_ring_t **rings, *ring;
6037     uint_t nrings;
6038     int rv = 0, i = 0, j;

6040     ASSERT((ring_type == MAC_RING_TYPE_RX &&
6041         mip->mi_rx_group_type == MAC_GROUP_TYPE_DYNAMIC) ||
6042         (ring_type == MAC_RING_TYPE_TX &&
6043         mip->mi_tx_group_type == MAC_GROUP_TYPE_DYNAMIC));

6045     /*
6046     * First find the rings to allocate to the group.
6047     */
6048     if (share != NULL) {
6049         /* get rings through ms_squery() */
6050         mip->mi_share_capab.ms_squery(share, ring_type, NULL, &nrings);
6051         ASSERT(nrings != 0);
6052         rings = kmem_alloc(nrings * sizeof (mac_ring_handle_t),
6053             KM_SLEEP);
6054         mip->mi_share_capab.ms_squery(share, ring_type,
6055             (mac_ring_handle_t *)rings, &nrings);
6056         for (i = 0; i < nrings; i++) {
6057             /*
6058             * If we have given this ring to a non-default
6059             * group, we need to check if we can get this
6060             * ring.
6061             */
6062             ring = rings[i];
6063             if (ring->mr_gh != (mac_group_handle_t)src_group ||
6064                 ring == (mac_ring_t *)mip->mi_default_tx_ring) {
6065                 if (mac_reclaim_ring_from_grp(mip, ring_type,
6066                     ring, rings, nrings) != 0) {
6067                     rv = ENOSPC;
6068                     goto bail;
6069                 }
6070             }
6071         }
6072     } else {
6073     }
6074 }
```

```

6073     /*
6074     * Pick one ring from default group.
6075     */
6076     /* for now pick the second ring which requires the first ring
6077     * at index 0 to stay in the default group, since it is the
6078     * ring which carries the multicast traffic.
6079     * We need a better way for a driver to indicate this,
6080     * for example a per-ring flag.
6081     */
6082     rings = kmem_alloc(ringcnt * sizeof (mac_ring_handle_t),
6083         KM_SLEEP);
6084     for (ring = src_group->mrg_rings; ring != NULL;
6085         ring = ring->mr_next) {
6086         if (ring_type == MAC_RING_TYPE_RX &&
6087             ring->mr_index == 0) {
6088             continue;
6089         }
6090         if (ring_type == MAC_RING_TYPE_TX &&
6091             ring == (mac_ring_t *)mip->mi_default_tx_ring) {
6092             continue;
6093         }
6094         rings[i++] = ring;
6095         if (i == ringcnt)
6096             break;
6097     }
6098     ASSERT(ring != NULL);
6099     nrings = i;
6100     /* Not enough rings as required */
6101     if (nrings != ringcnt) {
6102         rv = ENOSPC;
6103         goto bail;
6104     }
6105 }

6106 switch (ring_type) {
6107 case MAC_RING_TYPE_RX:
6108     if (src_group->mrg_cur_count - nrings < 1) {
6109         /* we ran out of rings */
6110         rv = ENOSPC;
6111         goto bail;
6112     }
6113
6114     /* move receive rings to new group */
6115     for (i = 0; i < nrings; i++) {
6116         rv = mac_group_mov_ring(mip, new_group, rings[i]);
6117         if (rv != 0) {
6118             /* move rings back on failure */
6119             for (j = 0; j < i; j++) {
6120                 (void) mac_group_mov_ring(mip,
6121                     src_group, rings[j]);
6122             }
6123             goto bail;
6124         }
6125     }
6126     break;
6127
6128 case MAC_RING_TYPE_TX: {
6129     mac_ring_t *tmp_ring;
6130
6131     /* move the TX rings to the new group */
6132     for (i = 0; i < nrings; i++) {
6133         /* get the desired ring */
6134         tmp_ring = mac_reserve_tx_ring(mip, rings[i]);
6135         if (tmp_ring == NULL) {
6136             rv = ENOSPC;
6137             goto bail;
6138         }
6139     }
6140 }
```

```

6139     }
6140     ASSERT(tmp_ring == rings[i]);
6141     rv = mac_group_mov_ring(mip, new_group, rings[i]);
6142     if (rv != 0) {
6143         /* cleanup on failure */
6144         for (j = 0; j < i; j++) {
6145             (void) mac_group_mov_ring(mip,
6146                                         MAC_DEFAULT_TX_GROUP(mip),
6147                                         rings[j]);
6148         }
6149         goto bail;
6150     }
6151 }
6152 break;
6153 }
6154 }

6155 /* add group to share */
6156 if (share != NULL)
6157     mip->mi_share_capab.ms_sadd(share, new_group->mrg_driver);

6158 bail:
6159 /* free temporary array of rings */
6160 kmem_free(rings, nrings * sizeof (mac_ring_handle_t));

6161 return (rv);
6162 }

6163 void
6164 mac_group_add_client(mac_group_t *grp, mac_client_impl_t *mcip)
6165 {
6166     mac_grp_client_t *mgcp;

6167     for (mgcp = grp->mrg_clients; mgcp != NULL; mgcp = mgcp->mfc_next)
6168         if (mgcp->mfc_client == mcip)
6169             break;
6170     }

6171 VERIFY(mgcp == NULL);

6172 mgcp = kmem_zalloc(sizeof (mac_grp_client_t), KM_SLEEP);
6173 mgcp->mfc_client = mcip;
6174 mgcp->mfc_next = grp->mrg_clients;
6175 grp->mrg_clients = mgcp;

6176 }

6177 void
6178 mac_group_remove_client(mac_group_t *grp, mac_client_impl_t *mcip)
6179 {
6180     mac_grp_client_t *mgcp, **pprev;

6181     for (pprev = &grp->mrg_clients, mgcp = *pprev; mgcp != NULL;
6182          pprev = &mgcp->mfc_next, mgcp = *pprev) {
6183         if (mgcp->mfc_client == mcip)
6184             break;
6185     }

6186     ASSERT(mgcp != NULL);

6187     *pprev = mgcp->mfc_next;
6188     kmem_free(mgcp, sizeof (mac_grp_client_t));
6189 }

6190 /*
6191 * mac_reserve_rx_group()
6192 */

```

```

6205 *
6206 * Finds an available group and exclusively reserves it for a client.
6207 * The group is chosen to suit the flow's resource controls (bandwidth and
6208 * fanout requirements) and the address type.
6209 * If the requestor is the primary MAC then return the group with the
6210 * largest number of rings, otherwise the default ring when available.
6211 */
6212 mac_group_t *
6213 mac_reserve_rx_group(mac_client_impl_t *mcip, uint8_t *mac_addr, boolean_t move)
6214 {
6215     mac_share_handle_t      share = mcip->mci_share;
6216     mac_impl_t              *mip = mcip->mci_mip;
6217     mac_group_t             *grp = NULL;
6218     int                      i;
6219     int                      err = 0;
6220     mac_address_t           *map;
6221     mac_resource_props_t    *mrp = MCIP_RESOURCE_PROPS(mcip);
6222     int                      nrings;
6223     int                      donor_grp_rcnt;
6224     boolean_t                need_exclgrp = B_FALSE;
6225     int                      need_rings = 0;
6226     mac_group_t              *candidate_grp = NULL;
6227     mac_client_impl_t        *gclient;
6228     mac_resource_props_t    *gmrp;
6229     mac_group_t              *donorgrp = NULL;
6230     boolean_t                rxhw = mrp->mrp_mask & MRP_RX_RINGS;
6231     boolean_t                unspec = mrp->mrp_mask & MRP_RXRINGS_UNSPEC;
6232     boolean_t                isprimary;

6234     ASSERT(MAC_PERIM_HELD((mac_handle_t)mip));
6235
6236     isprimary = mcip->mci_flemt->fe_type & FLOW_PRIMARY_MAC;

6238 /*
6239  * Check if a group already has this mac address (case of VLANs)
6240  * unless we are moving this MAC client from one group to another.
6241  */
6242     if (!move && (map = mac_find_macaddr(mip, mac_addr)) != NULL) {
6243         if (map->ma_group != NULL)
6244             return (map->ma_group);
6245     }
6246     if (mip->mi_rx_groups == NULL || mip->mi_rx_group_count == 0)
6247         return (NULL);
6248 /*
6249  * If exclusive open, return NULL which will enable the
6250  * caller to use the default group.
6251  */
6252     if (mcip->mci_state_flags & MCIS_EXCLUSIVE)
6253         return (NULL);

6255 /*
6256  * For dynamic groups default unspecified to 1 */
6257     if (rxhw && unspec &&
6258         mip->mi_rx_group_type == MAC_GROUP_TYPE_DYNAMIC) {
6259         mrp->mrp_nrxrings = 1;
6260     }
6261 /*
6262  * For static grouping we allow only specifying rings=0 and
6263  * unspecified
6264  */
6265     if (rxhw && mrp->mrp_nrxrings > 0 &&
6266         mip->mi_rx_group_type == MAC_GROUP_TYPE_STATIC) {
6267         return (NULL);
6268     }
6269     if (rxhw) {
6270         /*
6271          * We have explicitly asked for a group (with nrxrings,

```

```

6271     * if unspec).
6272     */
6273     if (unspec || mrp->mrp_nrxrings > 0) {
6274         need_exclgrp = B_TRUE;
6275         need_rings = mrp->mrp_nrxrings;
6276     } else if (mrp->mrp_nrxrings == 0) {
6277         /*
6278          * We have asked for a software group.
6279          */
6280         return (NULL);
6281     }
6282 } else if (isprimary && mip->mi_nactiveclients == 1 &&
6283     mip->mi_rx_group_type == MAC_GROUP_TYPE_DYNAMIC) {
6284     /*
6285      * If the primary is the only active client on this
6286      * mip and we have not asked for any rings, we give
6287      * it the default group so that the primary gets to
6288      * use all the rings.
6289      */
6290     return (NULL);
6291 }

6293 /* The group that can donate rings */
6294 donorgrp = mip->mi_rx_donor_grp;
6295
6296 /*
6297  * The number of rings that the default group can donate.
6298  * We need to leave at least one ring.
6299  */
6300 donor_grp_rcnt = donorgrp->mrg_cur_count - 1;

6302 /*
6303  * Try to exclusively reserve a RX group.
6304  *
6305  * For flows requiring HW_DEFAULT_RING (unicast flow of the primary
6306  * client), try to reserve the a non-default RX group and give
6307  * it all the rings from the donor group, except the default ring
6308  *
6309  * For flows requiring HW_RING (unicast flow of other clients), try
6310  * to reserve non-default RX group with the specified number of
6311  * rings, if available.
6312  *
6313  * For flows that have not asked for software or hardware ring,
6314  * try to reserve a non-default group with 1 ring, if available.
6315  */
6316 for (i = 1; i < mip->mi_rx_group_count; i++) {
6317     grp = &mip->mi_rx_groups[i];
6318
6319     DTRACE_PROBE3(rx_group_trying, char *, mip->mi_name,
6320                   int, grp->mrg_index, mac_group_state_t, grp->mrg_state);
6321
6322     /*
6323      * Check if this group could be a candidate group for
6324      * eviction if we need a group for this MAC client,
6325      * but there aren't any. A candidate group is one
6326      * that didn't ask for an exclusive group, but got
6327      * one and it has enough rings (combined with what
6328      * the donor group can donate) for the new MAC
6329      * client
6330      */
6331     if (grp->mrg_state >= MAC_GROUP_STATE_RESERVED) {
6332         /*
6333          * If the primary/donor group is not the default
6334          * group, don't bother looking for a candidate group.
6335          * If we don't have enough rings we will check
6336          * if the primary group can be vacated.
6337
6338
6339
6340
6341
6342
6343
6344
6345
6346
6347
6348
6349
6350
6351
6352
6353
6354
6355
6356
6357
6358
6359
6360
6361
6362
6363
6364
6365
6366
6367
6368
6369
6370
6371
6372
6373
6374
6375
6376
6377
6378
6379
6380
6381
6382
6383
6384
6385
6386
6387
6388
6389
6390
6391
6392
6393
6394
6395
6396
6397
6398
6399
6400
6401
6402
6403
6404
6405
6406
6407
6408
6409
6410
6411
6412
6413
6414
6415
6416
6417
6418
6419
6420
6421
6422
6423
6424
6425
6426
6427
6428
6429
6430
6431
6432
6433
6434
6435
6436
6437
6438
6439
6440
6441
6442
6443
6444
6445
6446
6447
6448
6449
6450
6451
6452
6453
6454
6455
6456
6457
6458
6459
6460
6461
6462
6463
6464
6465
6466
6467
6468
6469
6470
6471
6472
6473
6474
6475
6476
6477
6478
6479
6480
6481
6482
6483
6484
6485
6486
6487
6488
6489
6490
6491
6492
6493
6494
6495
6496
6497
6498
6499
6500
6501
6502
6503
6504
6505
6506
6507
6508
6509
6510
6511
6512
6513
6514
6515
6516
6517
6518
6519
6520
6521
6522
6523
6524
6525
6526
6527
6528
6529
6530
6531
6532
6533
6534
6535
6536
6537
6538
6539
6540
6541
6542
6543
6544
6545
6546
6547
6548
6549
6550
6551
6552
6553
6554
6555
6556
6557
6558
6559
6560
6561
6562
6563
6564
6565
6566
6567
6568
6569
6570
6571
6572
6573
6574
6575
6576
6577
6578
6579
6580
6581
6582
6583
6584
6585
6586
6587
6588
6589
6590
6591
6592
6593
6594
6595
6596
6597
6598
6599
6599
6600
6601
6602
6603
6604
6605
6606
6607
6608
6609
6609
6610
6611
6612
6613
6614
6615
6616
6617
6618
6619
6619
6620
6621
6622
6623
6624
6625
6626
6627
6628
6629
6629
6630
6631
6632
6633
6634
6635
6636
6637
6638
6639
6640
6641
6642
6643
6644
6645
6646
6647
6648
6649
6649
6650
6651
6652
6653
6654
6655
6656
6657
6658
6659
6659
6660
6661
6662
6663
6664
6665
6666
6667
6668
6669
6669
6670
6671
6672
6673
6674
6675
6676
6677
6678
6679
6679
6680
6681
6682
6683
6684
6685
6686
6687
6688
6689
6689
6690
6691
6692
6693
6694
6695
6696
6697
6698
6699
6699
6700
6701
6702
6703
6704
6705
6706
6707
6708
6709
6709
6710
6711
6712
6713
6714
6715
6716
6717
6718
6719
6719
6720
6721
6722
6723
6724
6725
6726
6727
6728
6729
6729
6730
6731
6732
6733
6734
6735
6736
6737
6738
6739
6739
6740
6741
6742
6743
6744
6745
6746
6747
6748
6749
6749
6750
6751
6752
6753
6754
6755
6756
6757
6758
6759
6759
6760
6761
6762
6763
6764
6765
6766
6767
6768
6769
6769
6770
6771
6772
6773
6774
6775
6776
6777
6778
6779
6779
6780
6781
6782
6783
6784
6785
6786
6787
6788
6789
6789
6790
6791
6792
6793
6794
6795
6796
6797
6798
6799
6799
6800
6801
6802
6803
6804
6805
6806
6807
6808
6809
6809
6810
6811
6812
6813
6814
6815
6816
6817
6818
6819
6819
6820
6821
6822
6823
6824
6825
6826
6827
6828
6829
6829
6830
6831
6832
6833
6834
6835
6836
6837
6838
6839
6840
6841
6842
6843
6844
6845
6846
6847
6848
6849
6849
6850
6851
6852
6853
6854
6855
6856
6857
6858
6859
6859
6860
6861
6862
6863
6864
6865
6866
6867
6868
6869
6869
6870
6871
6872
6873
6874
6875
6876
6877
6878
6879
6879
6880
6881
6882
6883
6884
6885
6886
6887
6888
6889
6889
6890
6891
6892
6893
6894
6895
6896
6897
6898
6899
6899
6900
6901
6902
6903
6904
6905
6906
6907
6908
6909
6909
6910
6911
6912
6913
6914
6915
6916
6917
6918
6919
6919
6920
6921
6922
6923
6924
6925
6926
6927
6928
6929
6929
6930
6931
6932
6933
6934
6935
6936
6937
6938
6939
6939
6940
6941
6942
6943
6944
6945
6946
6947
6948
6949
6949
6950
6951
6952
6953
6954
6955
6956
6957
6958
6959
6959
6960
6961
6962
6963
6964
6965
6966
6967
6968
6969
6969
6970
6971
6972
6973
6974
6975
6976
6977
6978
6979
6979
6980
6981
6982
6983
6984
6985
6986
6987
6988
6989
6989
6990
6991
6992
6993
6994
6995
6996
6997
6998
6999
6999
7000
7001
7002
7003
7004
7005
7006
7007
7008
7009
7009
7010
7011
7012
7013
7014
7015
7016
7017
7018
7019
7019
7020
7021
7022
7023
7024
7025
7026
7027
7028
7029
7029
7030
7031
7032
7033
7034
7035
7036
7037
7038
7039
7039
7040
7041
7042
7043
7044
7045
7046
7047
7048
7049
7049
7050
7051
7052
7053
7054
7055
7056
7057
7058
7059
7059
7060
7061
7062
7063
7064
7065
7066
7067
7068
7069
7069
7070
7071
7072
7073
7074
7075
7076
7077
7078
7079
7079
7080
7081
7082
7083
7084
7085
7086
7087
7088
7089
7089
7090
7091
7092
7093
7094
7095
7096
7097
7098
7099
7099
7100
7101
7102
7103
7104
7105
7106
7107
7108
7109
7109
7110
7111
7112
7113
7114
7115
7116
7117
7118
7119
7119
7120
7121
7122
7123
7124
7125
7126
7127
7128
7129
7129
7130
7131
7132
7133
7134
7135
7136
7137
7138
7139
7139
7140
7141
7142
7143
7144
7145
7146
7147
7148
7149
7149
7150
7151
7152
7153
7154
7155
7156
7157
7158
7159
7159
7160
7161
7162
7163
7164
7165
7166
7167
7168
7169
7169
7170
7171
7172
7173
7174
7175
7176
7177
7178
7179
7179
7180
7181
7182
7183
7184
7185
7186
7187
7188
7189
7189
7190
7191
7192
7193
7194
7195
7196
7197
7198
7199
7199
7200
7201
7202
7203
7204
7205
7206
7207
7208
7209
7209
7210
7211
7212
7213
7214
7215
7216
7217
7218
7219
7219
7220
7221
7222
7223
7224
7225
7226
7227
7228
7229
7229
7230
7231
7232
7233
7234
7235
7236
7237
7238
7239
7239
7240
7241
7242
7243
7244
7245
7246
7247
7248
7249
7249
7250
7251
7252
7253
7254
7255
7256
7257
7258
7259
7259
7260
7261
7262
7263
7264
7265
7266
7267
7268
7269
7269
7270
7271
7272
7273
7274
7275
7276
7277
7278
7279
7279
7280
7281
7282
7283
7284
7285
7286
7287
7288
7289
7289
7290
7291
7292
7293
7294
7295
7296
7297
7298
7299
7299
7300
7301
7302
7303
7304
7305
7306
7307
7308
7309
7309
7310
7311
7312
7313
7314
7315
7316
7317
7318
7319
7319
7320
7321
7322
7323
7324
7325
7326
7327
7328
7329
7329
7330
7331
7332
7333
7334
7335
7336
7337
7338
7339
7339
7340
7341
7342
7343
7344
7345
7346
7347
7348
7349
7349
7350
7351
7352
7353
7354
7355
7356
7357
7358
7359
7359
7360
7361
7362
7363
7364
7365
7366
7367
7368
7369
7369
7370
7371
7372
7373
7374
7375
7376
7377
7378
7379
7379
7380
7381
7382
7383
7384
7385
7386
7387
7388
7389
7389
7390
7391
7392
7393
7394
7395
7396
7397
7398
7399
7399
7400
7401
7402
7403
7404
7405
7406
7407
7408
7409
7409
7410
7411
7412
7413
7414
7415
7416
7417
7418
7419
7419
7420
7421
7422
7423
7424
7425
7426
7427
7428
7429
7429
7430
7431
7432
7433
7434
7435
7436
7437
7438
7439
7439
7440
7441
7442
7443
7444
7445
7446
7447
7448
7449
7449
7450
7451
7452
7453
7454
7455
7456
7457
7458
7459
7459
7460
7461
7462
7463
7464
7465
7466
7467
7468
7469
7469
7470
7471
7472
7473
7474
7475
7476
7477
7478
7479
7479
7480
7481
7482
7483
7484
7485
7486
7487
7488
7489
7489
7490
7491
7492
7493
7494
7495
7496
7497
7498
7499
7499
7500
7501
7502
7503
7504
7505
7506
7507
7508
7509
7509
7510
7511
7512
7513
7514
7515
7516
7517
7518
7519
7519
7520
7521
7522
7523
7524
7525
7526
7527
7528
7529
7529
7530
7531
7532
7533
7534
7535
7536
7537
7538
7539
7539
7540
7541
7542
7543
7544
7545
7546
7547
7548
7549
7549
7550
7551
7552
7553
7554
7555
7556
7557
7558
7559
7559
7560
7561
7562
7563
7564
7565
7566
7567
7568
7569
7569
7570
7571
7572
7573
7574
7575
7576
7577
7578
7579
7579
7580
7581
7582
7583
7584
7585
7586
7587
7588
7589
7589
7590
7591
7592
7593
7594
7595
7596
7597
7598
7599
7599
7600
7601
7602
7603
7604
7605
7606
7607
7608
7609
7609
7610
7611
7612
7613
7614
7615
7616
7617
7618
7619
7619
7620
7621
7622
7623
7624
7625
7626
7627
7628
7629
7629
7630
7631
7632
7633
7634
7635
7636
7637
7638
7639
7639
7640
7641
7642
7643
7644
7645
7646
7647
7648
7649
7649
7650
7651
7652
7653
7654
7655
7656
7657
7658
7659
7659
7660
7661
7662
7663
7664
7665
7666
7667
7668
7669
7669
7670
7671
7672
7673
7674
7675
7676
7677
7678
7679
7679
7680
7681
7682
7683
7684
7685
7686
7687
7688
7689
7689
7690
7691
7692
7693
7694
7695
7696
7697
7698
7699
7699
7700
7701
7702
7703
7704
7705
7706
7707
7708
7709
7709
7710
7711
7712
7713
7714
7715
7716
7717
7718
7719
7719
7720
7721
7722
7723
7724
7725
7726
7727
7728
7729
7729
7730
7731
7732
7733
7734
7735
7736
7737
7738
7739
7739
7740
7741
7742
7743
7744
7745
7746
7747
7748
7749
7749
7750
7751
7752
7753
7754
7755
7756
7757
7758
7759
7759
7760
7761
7762
7763
7764
7765
7766
7767
7768
7769
7769
7770
7771
7772
7773
7774
7775
7776
7777
7778
7779
7779
7780
7781
7782
7783
7784
7785
7786
7787
7788
7789
7789
7790
7791
7792
7793
7794
7795
7796
7797
7798
7799
7799
7800
7801
7802
7803
7804
7805
7806
7807
7808
7809
7809
7810
7811
7812
7813
7814
7815
7816
7817
7818
7819
7819
7820
7821
7822
7823
7824
7825
7826
7827
7828
7829
7829
7830
7831
7832
7833
7834
7835
7836
7837
7838
7839
7839
7840
7841
7842
7843
7844
7845
7846
7847
7848
7849
7849
7850
7851
7852
7853
7854
7855
7856
7857
7858
7859
7859
7860
7861
7862
7863
7864
7865
7866
7867
7868
7869
7869
7870
7871
7872
7873
7874
7875
7876
7877
7878
7879
7879
7880
7881
7882
7883
7884
7885
7886
7887
7888
7889
7889
7890
7891
7892
7893
7894
7895
7896
7897
7898
7899
7899
7900
7901
7902
7903
7904
7905
7906
7907
7908
7909
7909
7910
7911
7912
7913
7914
7915
7916
7917
7918
7919
7919
7920
7921
7922
7923
7924
7925
7926
7927
7928
7929
7929
7930
7931
7932
7933
7934
7935
7936
7937
7938
7939
7939
7940
7941
7942
7943
7944
7945
7946
7947
7948
7949
7949
7950
7951
7952
7953
7954
7955
7956
7957
7958
7959
7959
7960
7961
7962
7963
7964
7965
7966
7967
7968
7969
7969
7970
7971
7972
7973
7974
7975
7976
7977
7978
7979
7979
7980
7981
7982
7983
7984
7985
7986
7987
7988
7989
7989
7990
7991
7992
7993
7994
7995
7996
7997
7998
7999
7999
8000
8001
8002
8003
8004
8005
8006
8007
8008
8009
8009
8010
8011
8012
8013
8014
8015
8016
8017
8018
8019
8019
8020
8021
8022
8023
8024
8025
8026
8027
8028
8029
8029
8030
8031
8032
8033
8034
8035
8036
8037
8038
8039
8039
8040
8041
8042
8043
8044
8045
8046
8047
8048
8049
8049
8050
8051
8052
8053
8054
8055
8056
8057
8058
8059
8059
8060
8061
8062
8063
8064
8065
8066
8067
8068
8069
8069
8070
8071
8072
8073
8074
8075
8076
8077
8078
8079
8079
8080
8081
8082
8083
8084
8085
8086
8087
8088
8089
8089
8090
8091
8092
8093
8094
8095
8096
8097
8098
8099
8099
8100
8101
8102
8103
8104
8105
8106
8107
8108
8109
8109
8110
8111
8112
8113
8114
8115
8116
8117
8118
8119
8119
8120
8121
8122
8123
8124
8125
8126
8127
8128
8129
8129
8130
8131
8132
8133
8134
8135
8136
8137
8138
8139
8139
8140
8141
8142
8143
8144
8145
8146
8147
8148
8149
8149
8150
8151
8152
8153
8154
8155
8156
8157
8158
8159
8159
8160
8161
8162
8163
8164
8165
8166
8167
8168
8169
8169
8170
8171
8172
8173
8174
8175
8176
8177
8178
8179
8179
8180
8181
8182
8183
8184
8185
8186
8187
8188
8189
8189
8190
8191
8192
8193
8194
8195
8196
8197
8198
8199
8199
8200
8201
8202
8203
8204
8205
8206
8207
8208
8209
8209
8210
8211
8212
8213
8214
8215
8216
8217
8218
8219
8219
8220
8221
8222
8223
8224
8225
8226
8227
8228
8229
8229
8230
8231
8232
8233
8234
8235
8236
8237
8238
8239
8239
8240
8241
8242
8243
8244
8245
8246
8247
8248
8249
8249
8250
82
```

```

6404 DTRACE_PROBE3(rx_group_reserve_alloc_rings, char *,
6405     mip->mi_name, int, grp->mrg_index, int, err);
6406
6407 /* It's a dynamic group but the grouping operation
6408 * failed.
6409 */
6410 mac_stop_group(grp);
6411 }
6412 /* We didn't find an exclusive group for this MAC client */
6413 if (i >= mip->mi_rx_group_count) {
6414
6415     if (!need_exclgrp)
6416         return (NULL);
6417
6418 /*
6419 * If we found a candidate group then we switch the
6420 * MAC client from the candidate_group to the default
6421 * group and give the group to this MAC client. If
6422 * we didn't find a candidate_group, check if the
6423 * primary is in its own group and if it can make way
6424 * for this MAC client.
6425 */
6426 if (candidate_grp == NULL &&
6427     donorgrp != MAC_DEFAULT_RX_GROUP(mip) &&
6428     donorgrp->mrg_cur_count >= need_rings) {
6429     candidate_grp = donorgrp;
6430 }
6431 if (candidate_grp != NULL) {
6432     boolean_t      prim_grp = B_FALSE;
6433
6434 /*
6435 * Switch the MAC client from the candidate group
6436 * to the default group.. If this group was the
6437 * donor group, then after the switch we need
6438 * to update the donor group too.
6439 */
6440 grp = candidate_grp;
6441 gclient = MAC_GROUP_ONLY_CLIENT(grp);
6442 if (gclient == NULL)
6443     gclient = mac_get_grp_primary(grp);
6444 if (grp == mip->mi_rx_donor_grp)
6445     prim_grp = B_TRUE;
6446 if (mac_rx_switch_group(gclient, grp,
6447     MAC_DEFAULT_RX_GROUP(mip)) != 0) {
6448     return (NULL);
6449 }
6450 if (prim_grp) {
6451     mip->mi_rx_donor_grp =
6452         MAC_DEFAULT_RX_GROUP(mip);
6453     donorgrp = MAC_DEFAULT_RX_GROUP(mip);
6454 }
6455
6456 /*
6457 * Now give this group with the required rings
6458 * to this MAC client.
6459 */
6460 ASSERT(grp->mrg_state == MAC_GROUP_STATE_REGISTERED);
6461 if (mac_start_group(grp) != 0)
6462     return (NULL);
6463 if (mip->mi_rx_group_type != MAC_GROUP_TYPE_DYNAMIC)
6464     return (grp);
6465

```

```

6469     donor_grp_rcnt = donorgrp->mrg_cur_count - 1;
6470     ASSERT(grp->mrg_cur_count == 0);
6471     ASSERT(donor_grp_rcnt >= need_rings);
6472     err = i_mac_group_allocate_rings(mip, MAC_RING_TYPE_RX,
6473         donorgrp, grp, share, need_rings);
6474     if (err == 0) {
6475         /*
6476          * For a share i_mac_group_allocate_rings gets
6477          * the rings from the driver, let's populate
6478          * the property for the client now.
6479         */
6480         if (share != NULL) {
6481             mac_client_set_rings(
6482                 (mac_client_handle_t)mcip,
6483                 grp->mrg_cur_count, -1);
6484     }
6485     DTRACE_PROBE2(rx_group_reserved,
6486         char *, mip->mi_name, int, grp->mrg_index);
6487     return (grp);
6488 }
6489 DTRACE_PROBE3(rx_group_reserve_alloc_rings, char *,
6490     mip->mi_name, int, grp->mrg_index, int, err);
6491 mac_stop_group(grp);
6492 }
6493 return (NULL);
6494 }
6495 ASSERT(grp != NULL);
6496 DTRACE_PROBE2(rx_group_reserved,
6497     char *, mip->mi_name, int, grp->mrg_index);
6498 return (grp);
6499
6500 }
6501 /*
6502 * mac_rx_release_group()
6503 */
6504 /*
6505 * This is called when there are no clients left for the group.
6506 * The group is stopped and marked MAC_GROUP_STATE_REGISTERED,
6507 * and if it is a non default group, the shares are removed and
6508 * all rings are assigned back to default group.
6509 */
6510 void
6511 mac_release_rx_group(mac_client_impl_t *mcip, mac_group_t *group)
6512 {
6513     mac_impl_t          *mip = mcip->mci_mip;
6514     mac_ring_t           *ring;
6515
6516     ASSERT(group != MAC_DEFAULT_RX_GROUP(mip));
6517
6518     if (mip->mi_rx_donor_grp == group)
6519         mip->mi_rx_donor_grp = MAC_DEFAULT_RX_GROUP(mip);
6520
6521 /*
6522 * This is the case where there are no clients left. Any
6523 * SRS etc on this group have also be quiesced.
6524 */
6525 for (ring = group->mrg_rings; ring != NULL; ring = ring->mr_next) {
6526     if (ring->mr_classify_type == MAC_HW_CLASSIFIER) {
6527         ASSERT(group->mrg_state == MAC_GROUP_STATE_RESERVED);
6528         /*
6529          * Remove the SRS associated with the HW ring.
6530          * As a result, polling will be disabled.
6531         */
6532         ring->mr_srs = NULL;
6533     }
6534 ASSERT(group->mrg_state < MAC_GROUP_STATE_RESERVED ||
```

```

6535             ring->mr_state == MR_INUSE);
6536     if (ring->mr_state == MR_INUSE) {
6537         mac_stop_ring(ring);
6538         ring->mr_flag = 0;
6539     }
6540 }
6541 /* remove group from share */
6542 if (mcip->mci_share != NULL) {
6543     mip->mi_share_capab.ms_sremove(mcip->mci_share,
6544         group->mrg_driver);
6545 }
6546
6547 if (mip->mi_rx_group_type == MAC_GROUP_TYPE_DYNAMIC) {
6548     mac_ring_t *ring;
6549
6550     /*
6551      * Rings were dynamically allocated to group.
6552      * Move rings back to default group.
6553      */
6554     while ((ring = group->mrg_rings) != NULL) {
6555         (void) mac_group_mov_ring(mip, mip->mi_rx_donor_grp,
6556             ring);
6557     }
6558     mac_stop_group(group);
6559 }
6560 /*
6561  * Possible improvement: See if we can assign the group just released
6562  * to a another client of the mip
6563  */
6564
6565 }
6566 */
6567 * When we move the primary's mac address between groups, we need to also
6568 * take all the clients sharing the same mac address along with it (VLANs)
6569 * We remove the mac address for such clients from the group after quiescing
6570 * them. When we add the mac address we restart the client. Note that
6571 * the primary's mac address is removed from the group after all the
6572 * other clients sharing the address are removed. Similarly, the primary's
6573 * mac address is added before all the other client's mac address are
6574 * added. While grp is the group where the clients reside, tgrp is
6575 * the group where the addresses have to be added.
6576 */
6577 static void
6578 mac_rx_move_macaddr_prim(mac_client_impl_t *mcip, mac_group_t *grp,
6579     mac_group_t *tgrp, uint8_t *maddr, boolean_t add)
6580 {
6581     mac_impl_t          *mip = mcip->mci_mip;
6582     mac_grp_client_t    *mgcp = grp->mrg_clients;
6583     mac_client_impl_t   *gmcip;
6584     boolean_t            prim;
6585
6586     prim = (mcip->mci_state_flags & MCIS_UNICAST_HW) != 0;
6587
6588     /*
6589      * If the clients are in a non-default group, we just have to
6590      * walk the group's client list. If it is in the default group
6591      * (which will be shared by other clients as well, we need to
6592      * check if the unicast address matches mcip's unicast.
6593      */
6594     while (mgcp != NULL) {
6595         gmcip = mgcp->mgc_client;
6596         if (gmcip != mcip &&
6597             (grp != MAC_DEFAULT_RX_GROUP(mip) ||
6598              mcip->mci_unicast == gmcip->mci_unicast)) {
6599             if (!add) {
6600

```

```

6601             mac_rx_client_quiesce(
6602                 (mac_client_handle_t)gmcip);
6603             (void) mac_remove_macaddr(mcip->mci_unicast);
6604         } else {
6605             (void) mac_add_macaddr(mip, tgrp, maddr, prim);
6606             mac_rx_client_restart(
6607                 (mac_client_handle_t)gmcip);
6608         }
6609     }
6610     mgcp = mgcp->mgc_next;
6611 }
6612 }
6613 */
6614 /*
6615  * Move the MAC address from fgrp to tgrp. If this is the primary client,
6616  * we need to take any VLANs etc. together too.
6617  */
6618 static int
6619 mac_rx_move_macaddr(mac_client_impl_t *mcip, mac_group_t *fgrp,
6620     mac_group_t *tgrp)
6621 {
6622     mac_impl_t          *mip = mcip->mci_mip;
6623     uint8_t               maddr[MAXMACADDRLEN];
6624     int                  err = 0;
6625     boolean_t            prim;
6626     boolean_t            multicln = B_FALSE;
6627
6628     mac_rx_client_quiesce((mac_client_handle_t)mcip);
6629     ASSERT(mcip->mci_unicast != NULL);
6630     bcopy(mcip->mci_unicast->ma_addr, maddr, mcip->mci_unicast->ma_len);
6631
6632     prim = (mcip->mci_state_flags & MCIS_UNICAST_HW) != 0;
6633     if (mcip->mci_unicast->ma_nusers > 1) {
6634         mac_rx_move_macaddr_prim(mcip, fgrp, NULL, maddr, B_FALSE);
6635         multicln = B_TRUE;
6636     }
6637     ASSERT(mcip->mci_unicast->ma_nusers == 1);
6638     err = mac_remove_macaddr(mcip->mci_unicast);
6639     if (err != 0) {
6640         mac_rx_client_restart((mac_client_handle_t)mcip);
6641         if (multicln) {
6642             mac_rx_move_macaddr_prim(mcip, fgrp, fgrp, maddr,
6643                                         B_TRUE);
6644         }
6645     }
6646     return (err);
6647 }
6648 */
6649 /*
6650  * Program the H/W Classifier first, if this fails we need
6651  * not proceed with the other stuff.
6652  */
6653 if ((err = mac_add_macaddr(mip, tgrp, maddr, prim)) != 0) {
6654     /*
6655      * Revert back the H/W Classifier */
6656     if ((err = mac_add_macaddr(mip, fgrp, maddr, prim)) != 0) {
6657         /*
6658          * This should not fail now since it worked earlier,
6659          * should we panic?
6660          */
6661         cmn_err(CE_WARN,
6662                 "mac_rx_switch_group: switching %p back"
6663                 " to group %p failed!!", (void *)mcip,
6664                 (void *)fgrp);
6665     }
6666     mac_rx_client_restart((mac_client_handle_t)mcip);
6667     if (multicln) {
6668         mac_rx_move_macaddr_prim(mcip, fgrp, fgrp, maddr,

```

```

6667             B_TRUE);
6668     }
6669     return (err);
6670 }
6671 mcip->mci_unicast = mac_find_macaddr(mip, maddr);
6672 mac_rx_client_restart((mac_client_handle_t)mcip);
6673 if (multicnt)
6674     mac_rx_move_macaddr_prim(mcip, fgrp, tgrp, maddr, B_TRUE);
6675 return (err);
6676 }

6678 /*
6679 * Switch the MAC client from one group to another. This means we need
6680 * to remove the MAC address from the group, remove the MAC client,
6681 * teardown the SRSSs and revert the group state. Then, we add the client
6682 * to the destination group, set the SRSSs, and add the MAC address to the
6683 * group.
6684 */
6685 int
6686 mac_rx_switch_group(mac_client_impl_t *mcip, mac_group_t *fgrp,
6687                      mac_group_t *tgrp)
6688 {
6689     int                     err;
6690     mac_group_state_t      next_state;
6691     mac_client_impl_t      *group_only_mcip;
6692     mac_client_impl_t      *gmcip;
6693     mac_impl_t             *mip = mcip->mci_mip;
6694     mac_grp_client_t       *mgcp;

6696     ASSERT(fgrp == mcip->mci_flext->fe_rx_ring_group);

6698     if ((err = mac_rx_move_macaddr(mcip, fgrp, tgrp)) != 0)
6699         return (err);

6701    /*
6702     * The group might be reserved, but SRSSs may not be set up, e.g.
6703     * primary and its vlans using a reserved group.
6704     */
6705     if (fgrp->mrg_state == MAC_GROUP_STATE_RESERVED &
6706         MAC_GROUP_ONLY_CLIENT(fgrp) != NULL) {
6707         mac_rx_srs_group_teardown(mcip->mci_flext, B_TRUE);
6708     }
6709     if (fgrp != MAC_DEFAULT_RX_GROUP(mip)) {
6710         mgcp = fgrp->mrg_clients;
6711         while (mgcp != NULL) {
6712             gmcip = mgcp->mgc_client;
6713             mgcp = mgcp->mgc_next;
6714             mac_group_remove_client(fgrp, gmcip);
6715             mac_group_add_client(tgrp, gmcip);
6716             gmcip->mci_flext->fe_rx_ring_group = tgrp;
6717         }
6718         mac_release_rx_group(mcip, fgrp);
6719         ASSERT(MAC_GROUP_NO_CLIENT(fgrp));
6720         mac_set_group_state(fgrp, MAC_GROUP_STATE_REGISTERED);
6721     } else {
6722         mac_group_remove_client(fgrp, mcip);
6723         mac_group_add_client(tgrp, mcip);
6724         mcip->mci_flext->fe_rx_ring_group = tgrp;
6725         /*
6726          * If there are other clients (VLANs) sharing this address
6727          * we should be here only for the primary.
6728          */
6729         if (mcip->mci_unicast->ma_nusers > 1) {
6730             /*
6731              * We need to move all the clients that are using
6732              * this h/w address.

```

```

6733             */
6734     mgcp = fgrp->mrg_clients;
6735     while (mgcp != NULL) {
6736         gmcip = mgcp->mgc_client;
6737         mgcp = mgcp->mgc_next;
6738         if (mcip->mci_unicast == gmcip->mci_unicast) {
6739             mac_group_remove_client(fgrp, gmcip);
6740             mac_group_add_client(tgrp, gmcip);
6741             gmcip->mci_flext->fe_rx_ring_group =
6742                 tgrp;
6743         }
6744     }
6745 }
6746 /*
6747 * The default group will still take the multicast,
6748 * broadcast traffic etc., so it won't go to
6749 * MAC_GROUP_STATE_REGISTERED.
6750 */
6751 if (fgrp->mrg_state == MAC_GROUP_STATE_RESERVED)
6752     mac_rx_group_unmark(fgrp, MR_CONDEMNED);
6753     mac_set_group_state(fgrp, MAC_GROUP_STATE_SHARED);
6754 }
6755 next_state = mac_group_next_state(tgrp, &group_only_mcip,
6756                                   MAC_DEFAULT_RX_GROUP(mip), B_TRUE);
6757 mac_set_group_state(tgrp, next_state);
6758 /*
6759 * If the destination group is reserved, setup the SRSSs etc.
6760 */
6761 if (tgrp->mrg_state == MAC_GROUP_STATE_RESERVED) {
6762     mac_rx_srs_group_setup(mcip, mcip->mci_flext, SRST_LINK);
6763     mac_fanout_setup(mcip, mcip->mci_flext,
6764                       MCIP_RESOURCE_PROPS(mcip), mac_rx_deliver, mcip, NULL,
6765                       NULL);
6766     mac_rx_group_unmark(tgrp, MR_INCIPIENT);
6767 } else {
6768     mac_rx_switch_grp_to_sw(tgrp);
6769 }
6770 return (0);
6771 }

6773 /*
6774 * Reserves a TX group for the specified share. Invoked by mac_tx_srs_setup()
6775 * when a share was allocated to the client.
6776 */
6777 mac_group_t *
6778 mac_reserve_tx_group(mac_client_impl_t *mcip, boolean_t move)
6779 {
6780     mac_impl_t             *mip = mcip->mci_mip;
6781     mac_group_t             *grp = NULL;
6782     int                     rv;
6783     int                     i;
6784     int                     err;
6785     mac_group_t             *defgrp;
6786     mac_share_handle_t      mac_resource_props_t
6787     share = mcip->mci_share;
6788     *mrp = MCIP_RESOURCE_PROPS(mcip);
6789     int                     nrings;
6790     int                     defnrings;
6791     boolean_t               need_exclgrp = B_FALSE;
6792     mac_group_t             *candidate_grp = NULL;
6793     mac_client_impl_t        *gclient;
6794     mac_resource_props_t    unspec = mrp->mrp_mask & MRP_TXRINGS_UNSPEC;
6795     boolean_t               isprimary;
6796     boolean_t               txhw;
6797

```

```

6799     isprimary = mcip->mci_flemt->fe_type & FLOW_PRIMARY_MAC;
6800     /*
6801      * When we come here for a VLAN on the primary (dladm create-vlan),
6802      * we need to pair it along with the primary (to keep it consistent
6803      * with the RX side). So, we check if the primary is already assigned
6804      * to a group and return the group if so. The other way is also
6805      * true, i.e. the VLAN is already created and now we are plumbing
6806      * the primary.
6807     */
6808     if (!move && isprimary) {
6809         for (gclient = mip->mi_clients_list; gclient != NULL;
6810              gclient = gclient->mci_client_next) {
6811             if (gclient->mci_flemt->fe_type & FLOW_PRIMARY_MAC &&
6812                 gclient->mci_flemt->fe_tx_ring_group != NULL) {
6813                 return (gclient->mci_flemt->fe_tx_ring_group);
6814             }
6815         }
6816     }
6817
6818     if (mip->mi_tx_groups == NULL || mip->mi_tx_group_count == 0)
6819         return (NULL);
6820
6821     /* For dynamic groups, default unspec to 1 */
6822     if (txhw && unspec &&
6823         mip->mi_tx_group_type == MAC_GROUP_TYPE_DYNAMIC) {
6824         mrp->mrp_ntxrings = 1;
6825     }
6826     /*
6827      * For static grouping we allow only specifying rings=0 and
6828      * unspecified
6829     */
6830     if (txhw && mrp->mrp_ntxrings > 0 &&
6831         mip->mi_tx_group_type == MAC_GROUP_TYPE_STATIC) {
6832         return (NULL);
6833     }
6834
6835     if (txhw) {
6836         /*
6837          * We have explicitly asked for a group (with ntxrings,
6838          * * if unspec).
6839         */
6840         if (unspec || mrp->mrp_ntxrings > 0) {
6841             need_exclgrp = B_TRUE;
6842             need_rings = mrp->mrp_ntxrings;
6843         } else if (mrp->mrp_ntxrings == 0) {
6844             /*
6845              * We have asked for a software group.
6846              */
6847             return (NULL);
6848         }
6849     }
6850     defgrp = MAC_DEFAULT_TX_GROUP(mip);
6851     /*
6852      * The number of rings that the default group can donate.
6853      * We need to leave at least one ring - the default ring - in
6854      * this group.
6855     */
6856     defnrings = defgrp->mrg_cur_count - 1;
6857
6858     /*
6859      * Primary gets default group unless explicitly told not
6860      * to (i.e. rings > 0).
6861     */
6862     if (isprimary && !need_exclgrp)
6863         return (NULL);

```

```

6865     nrings = (mrp->mrp_mask & MRP_TX_RINGS) != 0 ? mrp->mrp_ntxrings : 1;
6866     for (i = 0; i < mip->mi_tx_group_count; i++) {
6867         grp = &mip->mi_tx_groups[i];
6868         if ((grp->mrg_state == MAC_GROUP_STATE_RESERVED) ||
6869             (grp->mrg_state == MAC_GROUP_STATE_UNINIT)) {
6870             /*
6871              * Select a candidate for replacement if we don't
6872              * get an exclusive group. A candidate group is one
6873              * that didn't ask for an exclusive group, but got
6874              * one and it has enough rings (combined with what
6875              * the default group can donate) for the new MAC
6876              * client.
6877             */
6878             if (grp->mrg_state == MAC_GROUP_STATE_RESERVED &&
6879                 candidate_grp == NULL) {
6880                 gclient = MAC_GROUP_ONLY_CLIENT(grp);
6881                 if (gclient == NULL)
6882                     gclient = mac_get_grp_primary(grp);
6883                 gmrp = MCIP_RESOURCE_PROPS(gclient);
6884                 if (gclient->mci_share == NULL &&
6885                     (gmrp->mrp_mask & MRP_TX_RINGS) == 0 &&
6886                     (unspec ||
6887                      (grp->mrg_cur_count + defnrings) >=
6888                      need_rings)) {
6889                     candidate_grp = grp;
6890                 }
6891             }
6892         }
6893         continue;
6894     }
6895     /*
6896      * If the default can't donate let's just walk and
6897      * see if someone can vacate a group, so that we have
6898      * enough rings for this.
6899     */
6900     if (mip->mi_tx_group_type != MAC_GROUP_TYPE_DYNAMIC ||
6901         nrings <= defnrings) {
6902         if (grp->mrg_state == MAC_GROUP_STATE_REGISTERED) {
6903             rv = mac_start_group(grp);
6904             ASSERT(rv == 0);
6905         }
6906     }
6907     break;
6908
6909     /*
6910      * The default group */
6911     if (i >= mip->mi_tx_group_count) {
6912         /*
6913          * If we need an exclusive group and have identified a
6914          * candidate group we switch the MAC client from the
6915          * candidate group to the default group and give the
6916          * candidate group to this client.
6917         */
6918         if (need_exclgrp && candidate_grp != NULL) {
6919             /*
6920              * Switch the MAC client from the candidate group
6921              * to the default group.
6922             */
6923             grp = candidate_grp;
6924             gclient = MAC_GROUP_ONLY_CLIENT(grp);
6925             if (gclient == NULL)
6926                 gclient = mac_get_grp_primary(grp);
6927             mac_tx_client_quiesce((mac_client_handle_t)gclient);
6928             mac_tx_switch_group(gclient, grp, defgrp);
6929             mac_tx_client_restart((mac_client_handle_t)gclient);
6930             /*

```

```

6931     * Give the candidate group with the specified number
6932     * of rings to this MAC client.
6933     */
6934     ASSERT(grp->mrg_state == MAC_GROUP_STATE_REGISTERED);
6935     rv = mac_start_group(grp);
6936     ASSERT(rv == 0);

6938
6939     if (mip->mi_tx_group_type != MAC_GROUP_TYPE_DYNAMIC)
6940         return (grp);

6941     ASSERT(grp->mrg_cur_count == 0);
6942     ASSERT(defgrp->mrg_cur_count > need_rings);

6944
6945     err = i_mac_group_allocate_rings(mip, MAC_RING_TYPE_TX,
6946         defgrp, grp, share, need_rings);
6947     if (err == 0) {
6948         /*
6949          * For a share i_mac_group_allocate_rings gets
6950          * the rings from the driver, let's populate
6951          * the property for the client now.
6952        */
6953         if (share != NULL) {
6954             mac_client_set_rings(
6955                 (mac_client_handle_t)mcip, -1,
6956                 grp->mrg_cur_count);
6957         }
6958         mip->mi_tx_group_free--;
6959     }
6960     DTRACE_PROBE3(tx_group_reserve_alloc_rings, char *,
6961         mip->mi_name, int, grp->mrg_index, int, err);
6962     mac_stop_group(grp);
6963 }
6964     return (NULL);
6965 }

6966 /**
6967  * We got an exclusive group, but it is not dynamic.
6968 */
6969 if (mip->mi_tx_group_type != MAC_GROUP_TYPE_DYNAMIC) {
6970     mip->mi_tx_group_free--;
6971     return (grp);
6972 }

6974
6975     rv = i_mac_group_allocate_rings(mip, MAC_RING_TYPE_TX, defgrp, grp,
6976         share, nrings);
6977     if (rv != 0) {
6978         DTRACE_PROBE3(tx_group_reserve_alloc_rings,
6979             char *, mip->mi_name, int, grp->mrg_index, int, rv);
6980         mac_stop_group(grp);
6981         return (NULL);
6982 }

6983 /**
6984  * For a share i_mac_group_allocate_rings gets the rings from the
6985  * driver, let's populate the property for the client now.
6986 */
6987 if (share != NULL) {
6988     mac_client_set_rings((mac_client_handle_t)mcip, -1,
6989         grp->mrg_cur_count);
6990 }
6991     mip->mi_tx_group_free--;
6992     return (grp);
6993 }

6994 void
6995 mac_release_tx_group(mac_client_impl_t *mcip, mac_group_t *grp)
6996 {

```

```

6997     mac_impl_t           *mip = mcip->mci_mip;
6998     mac_share_handle_t   share = mcip->mci_share;
6999     mac_ring_t            ring;
7000     mac_soft_ring_set_t  *srs = MCIP_TX_SRS(mcip);
7001     mac_group_t           defgrp;

7003     defgrp = MAC_DEFAULT_TX_GROUP(mip);
7004     if (srs != NULL) {
7005         if (srs->srs_soft_ring_count > 0) {
7006             for (ring = grp->mrg_rings; ring != NULL;
7007                 ring = ring->mr_next) {
7008                 ASSERT(mac_tx_srs_ring_present(srs, ring));
7009                 mac_tx_invoke_callbacks(mcip,
7010                     (mac_tx_cookie_t)
7011                     mac_tx_srs_get_soft_ring(srs, ring));
7012                 mac_tx_srs_del_ring(srs, ring);
7013             }
7014         } else {
7015             ASSERT(srs->srs_tx.st_arg2 != NULL);
7016             srs->srs_tx.st_arg2 = NULL;
7017             mac_srs_stat_delete(srs);
7018         }
7019     }
7020     if (share != NULL)
7021         mip->mi_share_capab.ms_sremove(share, grp->mrg_driver);

7023     /* move the ring back to the pool */
7024     if (mip->mi_tx_group_type == MAC_GROUP_TYPE_DYNAMIC) {
7025         while ((ring = grp->mrg_rings) != NULL)
7026             (void) mac_group_mov_ring(mip, defgrp, ring);
7027     }
7028     mac_stop_group(grp);
7029     mip->mi_tx_group_free++;

7032 /**
7033  * Disassociate a MAC client from a group, i.e go through the rings in the
7034  * group and delete all the soft rings tied to them.
7035 */
7036 static void
7037 mac_tx_dismantle_soft_rings(mac_group_t *fgrp, flow_entry_t *flent)
7038 {
7039     mac_client_impl_t      *mcip = flent->fe_mcip;
7040     mac_soft_ring_set_t    *tx_srs;
7041     mac_srs_tx_t           *tx;
7042     mac_ring_t              ring;

7044     tx_srs = flent->fe_tx_srs;
7045     tx = &tx_srs->srs_tx;

7047     /* Single ring case we haven't created any soft rings */
7048     if (tx->st_mode == SRS_TX_BW || tx->st_mode == SRS_TX_SERIALIZE ||
7049         tx->st_mode == SRS_TX_DEFAULT) {
7050         tx->st_arg2 = NULL;
7051         mac_srs_stat_delete(tx_srs);
7052     } /* Fanout case, where we have to dismantle the soft rings */
7053     else {
7054         for (ring = fgrp->mrg_rings; ring != NULL;
7055             ring = ring->mr_next) {
7056             ASSERT(mac_tx_srs_ring_present(tx_srs, ring));
7057             mac_tx_invoke_callbacks(mcip,
7058                 (mac_tx_cookie_t)mac_tx_srs_get_soft_ring(tx_srs,
7059                     ring));
7060             mac_tx_srs_del_ring(tx_srs, ring);
7061         }
7062     }
7063     ASSERT(tx->st_arg2 == NULL);

```



```

7195     mac_rx_deliver, mcip, NULL, NULL);
7196 }

7198 /*
7199  * This is a 1-time control path activity initiated by the client (IP).
7200  * The mac perimeter protects against other simultaneous control activities,
7201  * for example an ioctl that attempts to change the degree of fanout and
7202  * increase or decrease the number of softstrings associated with this Tx SRS.
7203 */
7204 static mac_tx_notify_cb_t *
7205 mac_client_tx_notify_add(mac_client_impl_t *mcip,
7206     mac_tx_notify_t notify, void *arg)
7207 {
7208     mac_cb_info_t *mcbi;
7209     mac_tx_notify_cb_t *mtnfp;

7211     ASSERT(MAC_PERIM_HELD((mac_handle_t)mcip->mci_mip));

7213     mtnfp = kmem_zalloc(sizeof (mac_tx_notify_cb_t), KM_SLEEP);
7214     mtnfp->mtnf_fn = notify;
7215     mtnfp->mtnf_arg = arg;
7216     mtnfp->mtnf_link.mcb_objp = mtnfp;
7217     mtnfp->mtnf_link.mcb_objsize = sizeof (mac_tx_notify_cb_t);
7218     mtnfp->mtnf_link.mcb_flags = MCB_TX_NOTIFY_CB_T;

7220     mcbi = &mcip->mci_tx_notify_cb_info;
7221     mutex_enter(mcbi->mcbi_lockp);
7222     mac_callback_add(mcbi, &mcip->mci_tx_notify_cb_list, &mtnfp->mtnf_link);
7223     mutex_exit(mcbi->mcbi_lockp);
7224     return (mtnfp);
7225 }

7227 static void
7228 mac_client_tx_notify_remove(mac_client_impl_t *mcip, mac_tx_notify_cb_t *mtnfp)
7229 {
7230     mac_cb_info_t *mcbi;
7231     mac_cb_t **cblist;

7233     ASSERT(MAC_PERIM_HELD((mac_handle_t)mcip->mci_mip));

7235     if (!mac_callback_find(&mcip->mci_tx_notify_cb_info,
7236         &mcip->mci_tx_notify_cb_list, &mtnfp->mtnf_link)) {
7237         cmn_err(CE_WARN,
7238             "mac_client_tx_notify_remove: callback not "
7239             "found, mcip 0x%p mtnfp 0x%p", (void *)mcip, (void *)mtnfp);
7240         return;
7241     }

7243     mcbi = &mcip->mci_tx_notify_cb_info;
7244     cblist = &mcip->mci_tx_notify_cb_list;
7245     mutex_enter(mcbi->mcbi_lockp);
7246     if (mac_callback_remove(mcbi, cblist, &mtnfp->mtnf_link))
7247         kmem_free(mtnfp, sizeof (mac_tx_notify_cb_t));
7248     else
7249         mac_callback_remove_wait(&mcip->mci_tx_notify_cb_info);
7250     mutex_exit(mcbi->mcbi_lockp);
7251 }

7253 /*
7254  * mac_client_tx_notify():
7255  * call to add and remove flow control callback routine.
7256 */
7257 mac_tx_notify_handle_t
7258 mac_client_tx_notify(mac_client_handle_t mch, mac_tx_notify_t callb_func,
7259     void *ptr)
7260 {

```

```

7261     mac_client_impl_t *mcip = (mac_client_impl_t *)mch;
7262     mac_tx_notify_cb_t *mtnfp = NULL;

7264     i_mac_perim_enter(mcip->mci_mip);

7266     if (callb_func != NULL) {
7267         /* Add a notify callback */
7268         mtnfp = mac_client_tx_notify_add(mcip, callb_func, ptr);
7269     } else {
7270         mac_client_tx_notify_remove(mcip, (mac_tx_notify_cb_t *)ptr);
7271     }
7272     i_mac_perim_exit(mcip->mci_mip);

7274     return ((mac_tx_notify_handle_t)mtnfp);
7275 }

7277 void
7278 mac_bridge_vectors(mac_bridge_tx_t txf, mac_bridge_rx_t rxf,
7279     mac_bridge_ref_t reff, mac_bridge_ls_t lsf)
7280 {
7281     mac_bridge_tx_cb = txf;
7282     mac_bridge_rx_cb = rxf;
7283     mac_bridge_ref_cb = reff;
7284     mac_bridge_ls_cb = lsf;
7285 }

7287 int
7288 mac_bridge_set(mac_handle_t mh, mac_handle_t link)
7289 {
7290     mac_impl_t *mip = (mac_impl_t *)mh;
7291     int retv;

7293     mutex_enter(&mip->mi_bridge_lock);
7294     if (mip->mi_bridge_link == NULL) {
7295         mip->mi_bridge_link = link;
7296         retv = 0;
7297     } else {
7298         retv = EBUSY;
7299     }
7300     mutex_exit(&mip->mi_bridge_lock);
7301     if (retv == 0) {
7302         mac_poll_state_change(mh, B_FALSE);
7303         mac_capab_update(mh);
7304     }
7305     return (retv);
7306 }

7308 /*
7309  * Disable bridging on the indicated link.
7310 */
7311 void
7312 mac_bridge_clear(mac_handle_t mh, mac_handle_t link)
7313 {
7314     mac_impl_t *mip = (mac_impl_t *)mh;

7316     mutex_enter(&mip->mi_bridge_lock);
7317     ASSERT(mip->mi_bridge_link == link);
7318     mip->mi_bridge_link = NULL;
7319     mutex_exit(&mip->mi_bridge_lock);
7320     mac_poll_state_change(mh, B_TRUE);
7321     mac_capab_update(mh);
7322 }

7324 void
7325 mac_no_active(mac_handle_t mh)
7326 {

```

```

7327     mac_impl_t *mip = (mac_impl_t *)mh;
7328
7329     i_mac_perim_enter(mip);
7330     mip->mi_state_flags |= MIS_NO_ACTIVE;
7331     i_mac_perim_exit(mip);
7332 }
7334 */
7335 * Walk the primary VLAN clients whenever the primary's rings property
7336 * changes and update the mac_resource_props_t for the VLAN's client.
7337 * We need to do this since we don't support setting these properties
7338 * on the primary's VLAN clients, but the VLAN clients have to
7339 * follow the primary w.r.t the rings property;
7340 */
7341 void
7342 mac_set_prim_vlan_rings(mac_impl_t *mip, mac_resource_props_t *mrp)
7343 {
7344     mac_client_impl_t      *vmcip;
7345     mac_resource_props_t   *vmrp;
7346
7347     for (vmcip = mip->mi_clients_list; vmcip != NULL;
7348          vmcip = vmcip->mci_client_next) {
7349         if (!!(vmcip->mci_client->fe_type & FLOW_PRIMARY_MAC) ||
7350             mac_client_vid(mac_client_handle_t)vmcip) ==
7351             VLAN_ID_NONE) {
7352             continue;
7353         }
7354         vmrp = MCIP_RESOURCE_PROPS(vmcip);
7355
7356         vmrp->mrp_nrtrings = mrp->mrp_nrtrings;
7357         if (mrp->mrp_mask & MRP_RX_RINGS)
7358             vmrp->mrp_mask |= MRP_RX_RINGS;
7359         else if (vmrp->mrp_mask & MRP_RX_RINGS)
7360             vmrp->mrp_mask &= ~MRP_RX_RINGS;
7361
7362         vmrp->mrp_ntxrings = mrp->mrp_ntxrings;
7363         if (mrp->mrp_mask & MRP_TX_RINGS)
7364             vmrp->mrp_mask |= MRP_TX_RINGS;
7365         else if (vmrp->mrp_mask & MRP_TX_RINGS)
7366             vmrp->mrp_mask &= ~MRP_TX_RINGS;
7367
7368         if (mrp->mrp_mask & MRP_RXRINGS_UNSPEC)
7369             vmrp->mrp_mask |= MRP_RXRINGS_UNSPEC;
7370         else
7371             vmrp->mrp_mask &= ~MRP_RXRINGS_UNSPEC;
7372
7373         if (mrp->mrp_mask & MRP_TXRINGS_UNSPEC)
7374             vmrp->mrp_mask |= MRP_TXRINGS_UNSPEC;
7375         else
7376             vmrp->mrp_mask &= ~MRP_TXRINGS_UNSPEC;
7377     }
7378 }
7379 */
7380 * We are adding or removing ring(s) from a group. The source for taking
7381 * rings is the default group. The destination for giving rings back is
7382 * the default group.
7383 */
7384 int
7385 mac_group_ring_modify(mac_client_impl_t *mcip, mac_group_t *group,
7386                         mac_group_t *defgrp)
7387 {
7388     mac_resource_props_t   *mrp = MCIP_RESOURCE_PROPS(mcip);
7389     uint_t                 modify;
7390     int                   count;
7391     mac_ring_t            *ring;

```

```

7393     mac_ring_t           *next;
7394     mac_impl_t           *mip = mcip->mci_mip;
7395     mac_ring_t           **rings;
7396     uint_t                ringcnt;
7397     int                  i = 0;
7398     boolean_t             rx_group = group->mrg_type == MAC_RING_TYPE_RX;
7399     int                  start;
7400     int                  end;
7401     mac_group_t           *tgrp;
7402     int                  j;
7403     int                  rv = 0;
7404
7405     /*
7406      * If we are asked for just a group, we give 1 ring, else
7407      * the specified number of rings.
7408      */
7409     if (rx_group) {
7410         ringcnt = (mrp->mrp_mask & MRP_RXRINGS_UNSPEC) ? 1:
7411                           mrp->mrp_nrxrings;
7412     } else {
7413         ringcnt = (mrp->mrp_mask & MRP_TXRINGS_UNSPEC) ? 1:
7414                           mrp->mrp_ntxrings;
7415     }
7416
7417     /* don't allow modifying rings for a share for now. */
7418     ASSERT(mcip->mci_share == NULL);
7419
7420     if (ringcnt == group->mrg_cur_count)
7421         return (0);
7422
7423     if (group->mrg_cur_count > ringcnt) {
7424         modify = group->mrg_cur_count - ringcnt;
7425         if (rx_group) {
7426             if (mip->mi_rx_donor_grp == group) {
7427                 ASSERT(mac_is_primary_client(mcip));
7428                 mip->mi_rx_donor_grp = defgrp;
7429             } else {
7430                 defgrp = mip->mi_rx_donor_grp;
7431             }
7432         }
7433         ring = group->mrg_rings;
7434         rings = kmem_alloc(modify * sizeof (mac_ring_handle_t),
7435                            KM_SLEEP);
7436         j = 0;
7437         for (count = 0; count < modify; count++) {
7438             next = ring->mr_next;
7439             rv = mac_group_mov_ring(mip, defgrp, ring);
7440             if (rv != 0) {
7441                 /* cleanup on failure */
7442                 for (j = 0; j < count; j++) {
7443                     (void) mac_group_mov_ring(mip, group,
7444                                         rings[j]);
7445                 }
7446                 break;
7447             }
7448             rings[j++] = ring;
7449             ring = next;
7450         }
7451         kmem_free(rings, modify * sizeof (mac_ring_handle_t));
7452         return (rv);
7453     }
7454     if (ringcnt >= MAX_RINGS_PER_GROUP)
7455         return (EINVAL);
7456
7457     modify = ringcnt - group->mrg_cur_count;

```

```

7459     if (rx_group) {
7460         if (group != mip->mi_rx_donor_grp)
7461             defgrp = mip->mi_rx_donor_grp;
7462         else
7463             /*
7464              * This is the donor group with all the remaining
7465              * rings. Default group now gets to be the donor
7466              */
7467             mip->mi_rx_donor_grp = defgrp;
7468         start = 1;
7469         end = mip->mi_rx_group_count;
7470     } else {
7471         start = 0;
7472         end = mip->mi_tx_group_count - 1;
7473     }
7474     /*
7475      * If the default doesn't have any rings, lets see if we can
7476      * take rings given to an h/w client that doesn't need it.
7477      * For now, we just see if there is any one client that can donate
7478      * all the required rings.
7479      */
7480     if (defgrp->mrg_cur_count < (modify + 1)) {
7481         for (i = start; i < end; i++) {
7482             if (rx_group) {
7483                 tgrp = &mip->mi_rx_groups[i];
7484                 if (tgrp == group || tgrp->mrg_state <
7485                     MAC_GROUP_STATE_RESERVED) {
7486                     continue;
7487                 }
7488                 mcip = MAC_GROUP_ONLY_CLIENT(tgrp);
7489                 if (mcip == NULL)
7490                     mcip = mac_get_grp_primary(tgrp);
7491                 ASSERT(mcip != NULL);
7492                 mrp = MCIP_RESOURCE_PROPS(mcip);
7493                 if ((mrp->mrp_mask & MRP_RX_RINGS) != 0)
7494                     continue;
7495                 if ((tgrp->mrg_cur_count +
7496                     defgrp->mrg_cur_count) < (modify + 1)) {
7497                     continue;
7498                 }
7499                 if (mac_rx_switch_group(mcip, tgrp,
7500                     defgrp) != 0) {
7501                     return (ENOSPC);
7502                 }
7503             } else {
7504                 tgrp = &mip->mi_tx_groups[i];
7505                 if (tgrp == group || tgrp->mrg_state <
7506                     MAC_GROUP_STATE_RESERVED) {
7507                     continue;
7508                 }
7509                 mcip = MAC_GROUP_ONLY_CLIENT(tgrp);
7510                 if (mcip == NULL)
7511                     mcip = mac_get_grp_primary(tgrp);
7512                 mrp = MCIP_RESOURCE_PROPS(mcip);
7513                 if ((mrp->mrp_mask & MRP_TX_RINGS) != 0)
7514                     continue;
7515                 if ((tgrp->mrg_cur_count +
7516                     defgrp->mrg_cur_count) < (modify + 1)) {
7517                     continue;
7518                 }
7519                 /* OK, we can switch this to s/w */
7520                 mac_tx_client_quiesce(
7521                     (mac_client_handle_t)mcip);
7522                 mac_tx_switch_group(mcip, tgrp, defgrp);
7523                 mac_tx_client_restart(
7524                     (mac_client_handle_t)mcip);

```

```

7525
7526
7527
7528
7529     }
7530     if (defgrp->mrg_cur_count < (modify + 1))
7531         return (ENOSPC);
7532     }
7533     if ((rv = i_mac_group_allocate_rings(mip, group->mrg_type, defgrp,
7534                                         group, mcip->mci_share, modify)) != 0) {
7535         return (rv);
7536     }
7537     return (0);
7538
7539     /*
7540      * Given the poolname in mac_resource_props, find the cpupart
7541      * that is associated with this pool. The cpupart will be used
7542      * later for finding the cpus to be bound to the networking threads.
7543      *
7544      * use_default is set B_TRUE if pools are enabled and pool_default
7545      * is returned. This avoids a 2nd lookup to set the poolname
7546      * for pool-effective.
7547      *
7548      * returns:
7549      *   NULL - pools are disabled or if the 'cpus' property is set.
7550      *   cpupart of pool_default - pools are enabled and the pool
7551      *     is not available or poolname is blank
7552      *   cpupart of named pool - pools are enabled and the pool
7553      *     is available.
7554      */
7555     cpupart_t *
7556     mac_pset_find(mac_resource_props_t *mrp, boolean_t *use_default)
7557     {
7558         pool_t          *pool;
7559         cpupart_t        *cpupart;
7560         *use_default = B_FALSE;
7561
7562         /* CPUs property is set */
7563         if (mrp->mrp_mask & MRP_CPU)
7564             return (NULL);
7565
7566         ASSERT(pool_lock_held());
7567
7568         /* Pools are disabled, no pset */
7569         if (pool_state == POOL_DISABLED)
7570             return (NULL);
7571
7572         /* Pools property is set */
7573         if (mrp->mrp_mask & MRP_POOL) {
7574             if ((pool = pool_lookup_pool_by_name(mrp->mrp_pool)) == NULL) {
7575                 /* Pool not found */
7576                 DTRACE_PROBE1(mac_pset_find_no_pool, char *,
7577                               mrp->mrp_pool);
7578                 *use_default = B_TRUE;
7579                 pool = pool_default;
7580             }
7581             /* Pools property is not set */
7582             } else {
7583                 *use_default = B_TRUE;
7584                 pool = pool_default;
7585             }
7586
7587             /* Find the CPU pset that corresponds to the pool */
7588             mutex_enter(&cpu_lock);
7589             if ((cpupart = cpupart_find(pool->pool_pset->pset_id)) == NULL) {
7590                 DTRACE_PROBE1(mac_find_pset_no_pset, psetid_t,
```

```
new/usr/src/uts/common/io/mac/mac.c

7591             pool->pool_pset->pset_id);
7592     }
7593     mutex_exit(&cpu_lock);
7595 }
7596 }

7598 void
7599 mac_set_pool_effective(boolean_t use_default, cpupart_t *cpupart,
7600     mac_resource_props_t *mrp, mac_resource_props_t *emrp)
7601 {
7602     ASSERT(pool_lock_held());
7604
7605     if (cpupart != NULL) {
7606         emrp->mrp_mask |= MRP_POOL;
7607         if (use_default) {
7608             (void) strcpy(emrp->mrp_pool,
7609                           "pool_default");
7610         } else {
7611             ASSERT(strlen(mrp->mrp_pool) != 0);
7612             (void) strcpy(emrp->mrp_pool,
7613                           mrp->mrp_pool);
7614         }
7615     } else {
7616         emrp->mrp_mask &= ~MRP_POOL;
7617         bzero(emrp->mrp_pool, MAXPATHLEN);
7618     }
7620 struct mac_pool_arg {
7621     char          mpa_poolname[MAXPATHLEN];
7622     pool_event_t   mpa_what;
7623 };
7625 /*ARGSUSED*/
7626 static uint_t
7627 mac_pool_link_update(mod_hash_key_t key, mod_hash_val_t *val, void *arg)
7628 {
7629     struct mac_pool_arg      *mpa = arg;
7630     mac_impl_t                *mip = (mac_impl_t *)val;
7631     mac_client_impl_t          *mcip;
7632     mac_resource_props_t       *mrp, *emrp;
7633     boolean_t                  pool_update = B_FALSE;
7634     boolean_t                  pool_clear = B_FALSE;
7635     boolean_t                  use_default = B_FALSE;
7636     cpupart_t                 *cpupart_t;
7638
7639     mrp = kmem_zalloc(sizeof (*mrp), KM_SLEEP);
7640     i_mac_perim_enter(mip);
7641     for (mcip = mip->mi_clients_list; mcip != NULL;
7642          mcip = mcip->mci_client_next) {
7643         pool_update = B_FALSE;
7644         pool_clear = B_FALSE;
7645         use_default = B_FALSE;
7646         mac_client_get_resources((mac_client_handle_t)mcip, mrp);
7647         emrp = MCIP_EFFECTIVE_PROPS(mcip);

7648     /*
7649     * When pools are enabled
7650     */
7651     if ((mpa->mpa_what == POOL_E_ENABLE) &&
7652         ((mrp->mrp_mask & MRP_CPUS) == 0)) {
7653         mrp->mrp_mask |= MRP_POOL;
7654         pool_update = B_TRUE;
7655     }
7656 }
```

```

7723     }
7724     i_mac_perim_exit(mip);
7725     kmem_free(mrp, sizeof (*mrp));
7726     return (MH_WALK_CONTINUE);
7727 }

7728 static void
7729 mac_pool_update(void *arg)
7730 {
7731     mod_hash_walk(i_mac_impl_hash, mac_pool_link_update, arg);
7732     kmem_free(arg, sizeof (struct mac_pool_arg));
7733 }
7734 */

7735 /* Callback function to be executed when a noteworthy pool event
7736 * takes place.
7737 */
7738 /* ARGSUSED */
7739 static void
7740 mac_pool_event_cb(pool_event_t what, poolid_t id, void *arg)
7741 {
7742     pool_t          *pool;
7743     char            *poolname = NULL;
7744     struct mac_pool_arg *mpa;
7745
7746     pool_lock();
7747     mpa = kmem_zalloc(sizeof (struct mac_pool_arg), KM_SLEEP);
7748
7749     switch (what) {
7750     case POOL_E_ENABLE:
7751     case POOL_E_DISABLE:
7752         break;
7753
7754     case POOL_E_CHANGE:
7755         pool = pool_lookup_pool_by_id(id);
7756         if (pool == NULL) {
7757             kmem_free(mpa, sizeof (struct mac_pool_arg));
7758             pool_unlock();
7759             return;
7760         }
7761         pool_get_name(pool, &poolname);
7762         (void) strlcpy(mpa->mpa_poolname, poolname,
7763                       sizeof (mpa->mpa_poolname));
7764         break;
7765
7766     default:
7767         kmem_free(mpa, sizeof (struct mac_pool_arg));
7768         pool_unlock();
7769         return;
7770     }
7771     pool_unlock();
7772
7773     mpa->mpa_what = what;
7774
7775     mac_pool_update(mpa);
7776
7777 }

7778 */
7779 * Set effective rings property. This could be called from datapath_setup/
7780 * datapath_teardown or set-linkprop.
7781 * If the group is reserved we just go ahead and set the effective rings.
7782 * Additionally, for TX this could mean the default group has lost/gained
7783 * some rings, so if the default group is reserved, we need to adjust the
7784 * effective rings for the default group clients. For RX, if we are working
7785 * with the non-default group, we just need * to reset the effective props
7786 * for the default group clients.
7787
7788

```

```

7789 */
7790 void
7791 mac_set_rings_effective(mac_client_impl_t *mcip)
7792 {
7793     mac_impl_t          *mip = mcip->mci_mip;
7794     mac_group_t          *grp;
7795     mac_group_t          *defgrp;
7796     flow_entry_t          *flent = mcip->mci_flent;
7797     mac_resource_props_t *emrp = MCIP_EFFECTIVE_PROPS(mcip);
7798     mac_grp_client_t      *mgcp;
7799     mac_client_impl_t      *gmcip;

7800     grp = flent->fe_rx_ring_group;
7801     if (grp != NULL) {
7802         defgrp = MAC_DEFAULT_RX_GROUP(mip);
7803         /*
7804          * If we have reserved a group, set the effective rings
7805          * to the ring count in the group.
7806          */
7807
7808         if (grp->mrg_state == MAC_GROUP_STATE_RESERVED) {
7809             emrp->mrp_mask |= MRP_RX_RINGS;
7810             emrp->mrp_nrxrings = grp->mrg_cur_count;
7811         }
7812
7813         /*
7814          * We go through the clients in the shared group and
7815          * reset the effective properties. It is possible this
7816          * might have already been done for some client (i.e.
7817          * if some client is being moved to a group that is
7818          * already shared). The case where the default group is
7819          * RESERVED is taken care of above (note in the RX side if
7820          * there is a non-default group, the default group is always
7821          * SHARED).
7822          */
7823         if (grp != defgrp || grp->mrg_state == MAC_GROUP_STATE_SHARED) {
7824             if (grp->mrg_state == MAC_GROUP_STATE_SHARED)
7825                 mgcp = grp->mrg_clients;
7826             else
7827                 mgcp = defgrp->mrg_clients;
7828             while (mgcp != NULL) {
7829                 gmcip = mgcp->mgc_client;
7830                 emrp = MCIP_EFFECTIVE_PROPS(gmcip);
7831                 if (emrp->mrp_mask & MRP_RX_RINGS) {
7832                     emrp->mrp_mask &= ~MRP_RX_RINGS;
7833                     emrp->mrp_nrxrings = 0;
7834                 }
7835             }
7836         }
7837     }
7838 }

7839 /* Now the TX side */
7840 grp = flent->fe_tx_ring_group;
7841 if (grp != NULL) {
7842     defgrp = MAC_DEFAULT_TX_GROUP(mip);
7843
7844     if (grp->mrg_state == MAC_GROUP_STATE_RESERVED) {
7845         emrp->mrp_mask |= MRP_TX_RINGS;
7846         emrp->mrp_ntxrings = grp->mrg_cur_count;
7847     } else if (grp->mrg_state == MAC_GROUP_STATE_SHARED) {
7848         mgcp = grp->mrg_clients;
7849         while (mgcp != NULL) {
7850             gmcip = mgcp->mgc_client;
7851             emrp = MCIP_EFFECTIVE_PROPS(gmcip);
7852             if (emrp->mrp_mask & MRP_TX_RINGS) {
7853                 emrp->mrp_mask &= ~MRP_TX_RINGS;
7854             }
7855         }
7856     }
7857 }
7858
7859

```

```

7855                     emrp->mrp_ntxrings = 0;
7856
7857             mgcp = mgcp->mgc_next;
7858         }
7859     }
7860
7861     /*
7862      * If the group is not the default group and the default
7863      * group is reserved, the ring count in the default group
7864      * might have changed, update it.
7865    */
7866    if (grp != defgrp &&
7867        defgrp->mrg_state == MAC_GROUP_STATE_RESERVED) {
7868        gmcip = MAC_GROUP_ONLY_CLIENT(defgrp);
7869        emrp = MCIP_EFFECTIVE_PROPS(gmcip);
7870        emrp->mrp_ntxrings = defgrp->mrg_cur_count;
7871    }
7872
7873    emrp = MCIP_EFFECTIVE_PROPS(mcip);
7874 }

7875 /*
7876  * Check if the primary is in the default group. If so, see if we
7877  * can give it a an exclusive group now that another client is
7878  * being configured. We take the primary out of the default group
7879  * because the multicast/broadcast packets for the all the clients
7880  * will land in the default ring in the default group which means
7881  * any client in the default group, even if it is the only one in
7882  * the group, will lose exclusive access to the rings, hence
7883  * polling.
7884 */
7885 mac_client_impl_t *
7886 mac_check_primary_relocation(mac_client_impl_t *mcip, boolean_t rxhw)
7887 {
7888     mac_impl_t          *mip = mcip->mci_mip;
7889     mac_group_t          *defgrp = MAC_DEFAULT_RX_GROUP(mip);
7890     flow_entry_t          *fleent = mcip->mci_fleent;
7891     mac_resource_props_t  *mrp = MCIP_RESOURCE_PROPS(mcip);
7892     uint8_t                *mac_addr;
7893     mac_group_t          *ngrp;
7894

7895 /*
7896  * Check if the primary is in the default group, if not
7897  * or if it is explicitly configured to be in the default
7898  * group OR set the RX rings property, return.
7899 */
7900 if (fleent->fe_rx_ring_group != defgrp || mrp->mrp_mask & MRP_RX_RINGS)
7901     return (NULL);

7902 /*
7903  * If the new client needs an exclusive group and we
7904  * don't have another for the primary, return.
7905 */
7906 if (rxhw && mip->mi_rxhwclnt_avail < 2)
7907     return (NULL);

7908 mac_addr = fleent->fe_flow_desc.fd_dst_mac;
7909 /*
7910  * We call this when we are setting up the datapath for
7911  * the first non-primary.
7912 */
7913 ASSERT(mip->mi_nactiveclients == 2);
7914 /*
7915  * OK, now we have the primary that needs to be relocated.
7916 */
7917 ngrp = mac_reserve_rx_group(mcip, mac_addr, B_TRUE);
7918
7919
7920

```

```

7921     if (ngrp == NULL)
7922         return (NULL);
7923     if (mac_rx_switch_group(mcip, defgrp, ngrp) != 0) {
7924         mac_stop_group(ngrp);
7925         return (NULL);
7926     }
7927     return (mcip);
7928 }
```

new/usr/src/uts/common/io/mac/mac\_client.c

1

```
*****  
148147 Thu Feb 20 18:59:06 2014  
new/usr/src/uts/common/io/mac/mac_client.c  
2553 mac address should be a dladm link property  
*****  
_____ unchanged_portion_omitted _____  
  
2344 /*  
2345 * Add a new unicast address to the MAC client.  
2346 *  
2347 * The MAC address can be specified either by value, or the MAC client  
2348 * can specify that it wants to use the primary MAC address of the  
2349 * underlying MAC. See the introductory comments at the beginning  
2350 * of this file for more information on primary MAC addresses.  
2350 * of this file for more more information on primary MAC addresses.  
2351 *  
2352 * Note also the tuple (MAC address, VID) must be unique  
2353 * for the MAC clients defined on top of the same underlying MAC  
2354 * instance, unless the MAC_UNICAST_NODUPCHECK is specified.  
2355 *  
2356 * In no case can a client use the PVID for the MAC, if the MAC has one set.  
2357 */  
2358 int  
2359 i_mac_unicast_add(mac_client_handle_t mch, uint8_t *mac_addr, uint16_t flags,  
2360     mac_unicast_handle_t *mah, uint16_t vid, mac_diag_t *diag)  
2361 {  
2362     mac_client_impl_t      *mcip = (mac_client_impl_t *)mch;  
2363     mac_impl_t             *mip = mcip->mci_mip;  
2364     int                     err;  
2365     uint_t                 mac_len = mip->mi_type->mt_addr_length;  
2366     boolean_t              check_dups = !(flags & MAC_UNICAST_NODUPCHECK);  
2367     boolean_t              fastpath_disabled = B_FALSE;  
2368     boolean_t              is_primary = (flags & MAC_UNICAST_PRIMARY);  
2369     boolean_t              is_unicast_hw = (flags & MAC_UNICAST_HW);  
2370     mac_resource_props_t   *mrp;  
2371     boolean_t              passive_client = B_FALSE;  
2372     mac_unicast_impl_t    *muip;  
2373     boolean_t              is_vnic_primary =  
2374         (flags & MAC_UNICAST_VNIC_PRIMARY);  
  
2375     /* when VID is non-zero, the underlying MAC can not be VNIC */  
2376     ASSERT(!((mip->mi_state_flags & MIS_IS_VNIC) && (vid != 0)));  
  
2377     /*  
2378      * Can't unicast add if the client asked only for minimal datapath  
2379      * setup.  
2380      */  
2381     if (mcip->mci_state_flags & MCIS_NO_UNICAST_ADDR)  
2382         return (ENOTSUP);  
  
2383     /*  
2384      * Check for an attempted use of the current Port VLAN ID, if enabled.  
2385      * No client may use it.  
2386      */  
2387     if (mip->mi_pvid != 0 && vid == mip->mi_pvid)  
2388         return (EBUSY);  
  
2389     /*  
2390      * Check whether it's the primary client and flag it.  
2391      */  
2392     if (!(mcip->mci_state_flags & MCIS_IS_VNIC) && is_primary && vid == 0)  
2393         mcip->mci_flags |= MAC_CLIENT_FLAGS_PRIMARY;  
  
2394     /*  
2395      * is_vnic_primary is true when we come here as a VLAN VNIC  
2396      * which uses the primary mac client's address but with a non-zero  
2397      */
```

new/usr/src/uts/common/io/mac/mac\_client.c

2

```
2402             * VID. In this case the MAC address is not specified by an upper  
2403             * MAC client.  
2404             */  
2405             if ((mcip->mci_state_flags & MCIS_IS_VNIC) && is_primary &&  
2406                 !is_vnic_primary) {  
2407                 /*  
2408                  * The address is being set by the upper MAC client  
2409                  * of a VNIC. The MAC address was already set by the  
2410                  * VNIC driver during VNIC creation.  
2411                  *  
2412                  * Note: a VNIC has only one MAC address. We return  
2413                  * the MAC unicast address handle of the lower MAC client  
2414                  * corresponding to the VNIC. We allocate a new entry  
2415                  * which is flagged appropriately, so that mac_unicast_remove()  
2416                  * doesn't attempt to free the original entry that  
2417                  * was allocated by the VNIC driver.  
2418                  */  
2419             ASSERT(mcip->mci_unicast != NULL);  
  
2420             /* Check for VLAN flags, if present */  
2421             if ((flags & MAC_UNICAST_TAG_DISABLE) != 0)  
2422                 mcip->mci_state_flags |= MCIS_TAG_DISABLE;  
  
2423             if ((flags & MAC_UNICAST_STRIP_DISABLE) != 0)  
2424                 mcip->mci_state_flags |= MCIS_STRIP_DISABLE;  
  
2425             if ((flags & MAC_UNICAST_DISABLE_TX_VID_CHECK) != 0)  
2426                 mcip->mci_state_flags |= MCIS_DISABLE_TX_VID_CHECK;  
  
2427             /*  
2428              * Ensure that the primary unicast address of the VNIC  
2429              * is added only once unless we have the  
2430              * MAC_CLIENT_FLAGS_MULTI_PRIMARY set (and this is not  
2431              * a passive MAC client).  
2432              */  
2433             if ((mcip->mci_flags & MAC_CLIENT_FLAGS_VNIC_PRIMARY) != 0) {  
2434                 if ((mcip->mci_flags &  
2435                     MAC_CLIENT_FLAGS_MULTI_PRIMARY) == 0 ||  
2436                     (mcip->mci_flags &  
2437                     MAC_CLIENT_FLAGS_PASSIVE_PRIMARY) != 0) {  
2438                     return (EBUSY);  
2439                 }  
2440                 mcip->mci_flags |= MAC_CLIENT_FLAGS_PASSIVE_PRIMARY;  
2441                 passive_client = B_TRUE;  
2442             }  
2443             mcip->mci_flags |= MAC_CLIENT_FLAGS_VNIC_PRIMARY;  
2444             /*  
2445              * Create a handle for vid 0.  
2446              */  
2447             ASSERT(vid == 0);  
2448             muip = kmem_zalloc(sizeof (mac_unicast_impl_t), KM_SLEEP);  
2449             muip->mui_vid = vid;  
2450             *mah = (mac_unicast_handle_t)muip;  
2451             /*  
2452              * This will be used by the caller to defer setting the  
2453              * rx functions.  
2454              */  
2455             if (passive_client)  
2456                 return (EAGAIN);  
2457             return (0);  
2458         }  
2459         /* primary MAC clients cannot be opened on top of anchor VNICs */  
2460         if ((is_vnic_primary || is_primary) &&
```

new/usr/src/uts/common/io/mac/mac\_client.c

```

2468     i_mac_capab_get((mac_handle_t)mip, MAC_CAPAB_ANCHOR_VNIC, NULL));
2469     return (ENXIO);
2470 }
2471
2472 /*
2473  * If this is a VNIC/VLAN, disable softmac fast-path.
2474  */
2475 if (mcip->mci_state_flags & MCIS_IS_VNIC) {
2476     err = mac_fastpath_disable((mac_handle_t)mip);
2477     if (err != 0)
2478         return (err);
2479     fastpath_disabled = B_TRUE;
2480 }
2481
2482 /*
2483  * Return EBUSY if:
2484  * - there is an exclusively active mac client exists.
2485  * - this is an exclusive active mac client but
2486  *   a. there is already active mac clients exist, or
2487  *   b. fastpath streams are already plumbed on this legacy device
2488  * - the mac creator has disallowed active mac clients.
2489 */
2490 if (mip->mi_state_flags & (MIS_EXCLUSIVE|MIS_NO_ACTIVE)) {
2491     if (fastpath_disabled)
2492         mac_fastpath_enable((mac_handle_t)mip);
2493     return (EBUSY);
2494 }
2495
2496 if (mcip->mci_state_flags & MCIS_EXCLUSIVE) {
2497     ASSERT(!fastpath_disabled);
2498     if (mip->mi_nactiveclients != 0)
2499         return (EBUSY);
2500
2501     if ((mip->mi_state_flags & MIS_LEGACY) &&
2502         !(mip->mi_capab_legacy.mi_active_set(mip->mi_driver))) {
2503         return (EBUSY);
2504     }
2505     mip->mi_state_flags |= MIS_EXCLUSIVE;
2506 }
2507
2508 mrp = kmem_zalloc(sizeof (*mrp), KM_SLEEP);
2509 if (is_primary && !(mcip->mci_state_flags & (MCIS_IS_VNIC |
2510     MCIS_IS_AGGR_PORT))) {
2511 /*
2512  * Apply the property cached in the mac_impl_t to the primary
2513  * mac Client. If the mac client is a VNIC or an aggregation
2514  * port, its property should be set in the mcip when the
2515  * VNIC/aggr was created.
2516  */
2517     mac_get_resources((mac_handle_t)mip, mrp);
2518     (void) mac_client_set_resources(mch, mrp);
2519 } else if (mcip->mci_state_flags & MCIS_IS_VNIC) {
2520 /*
2521  * This is a primary VLAN client, we don't support
2522  * specifying rings property for this as it inherits the
2523  * rings property from its MAC.
2524  */
2525     if (is_vnic_primary) {
2526         mac_resource_props_t *vmrp;
2527
2528         vmrp = MCIP_RESOURCE_PROPS(mcip);
2529         if (vmrp->mrp_mask & MRP_RX_RINGS ||
2530             vmrp->mrp_mask & MRP_TX_RINGS) {
2531             if (fastpath_disabled)
2532                 mac_fastpath_enable((mac_handle_t)mip);
2533             kmem_free(mrp, sizeof (*mrp));
2534         }
2535     }
2536 }

```

new/usr/src/uts/common/io/mac/mac\_client.c

```

2534         return (ENOTSUP);
2535     }
2536     /*
2537      * Additionally we also need to inherit any
2538      * rings property from the MAC.
2539      */
2540     mac_get_resources((mac_handle_t)mip, mrp);
2541     if (mrp->mrp_mask & MRP_RX_RINGS) {
2542         vmrp->mrp_mask |= MRP_RX_RINGS;
2543         vmrp->mrp_nrxrings = mrp->mrp_nrxrings;
2544     }
2545     if (mrp->mrp_mask & MRP_TX_RINGS) {
2546         vmrp->mrp_mask |= MRP_TX_RINGS;
2547         vmrp->mrp_ntxrings = mrp->mrp_ntxrings;
2548     }
2549 }
2550 bcopy(MCIP_RESOURCE_PROPS(mcip), mrp, sizeof (*mrp));
2551 }

2553 muip = kmalloc(sizeof (mac_unicast_impl_t), KM_SLEEP);
2554 muip->mui_vid = vid;

2556 if (is_primary || is_vnic_primary) {
2557     mac_addr = mip->mi_addr;
2558 } else {

2560     /*
2561      * Verify the validity of the specified MAC addresses value.
2562      */
2563     if (!mac_unicst_verify((mac_handle_t)mip, mac_addr, mac_len)) {
2564         *diag = MAC_DIAG_MACADDR_INVALID;
2565         err = EINVAL;
2566         goto bail_out;
2567     }

2569     /*
2570      * Make sure that the specified MAC address is different
2571      * than the unicast MAC address of the underlying NIC.
2572      */
2573     if (check_dups && bcmp(mip->mi_addr, mac_addr, mac_len) == 0) {
2574         *diag = MAC_DIAG_MACADDR_NIC;
2575         err = EINVAL;
2576         goto bail_out;
2577     }
2578 }

2580 /*
2581  * Set the flags here so that if this is a passive client, we
2582  * can return and set it when we call mac_client_datapath_setup
2583  * when this becomes the active client. If we defer to using these
2584  * flags to mac_client_datapath_setup, then for a passive client,
2585  * we'd have to store the flags somewhere (probably fe_flags)
2586  * and then use it.
2587 */
2588 if (!MCIP_DATAPATH_SETUP(mcip)) {
2589     if (is_unicast_hw) {
2590         /*
2591          * The client requires a hardware MAC address slot
2592          * for that unicast address. Since we support only
2593          * one unicast MAC address per client, flag the
2594          * MAC client itself.
2595          */
2596         mcip->mci_state_flags |= MCIS_UNICAST_HW;
2597     }
2599 /* Check for VLAN flags, if present */

```

```

2600     if ((flags & MAC_UNICAST_TAG_DISABLE) != 0)
2601         mcip->mci_state_flags |= MCIS_TAG_DISABLE;
2602
2603     if ((flags & MAC_UNICAST_STRIP_DISABLE) != 0)
2604         mcip->mci_state_flags |= MCIS_STRIP_DISABLE;
2605
2606     if ((flags & MAC_UNICAST_DISABLE_TX_VID_CHECK) != 0)
2607         mcip->mci_state_flags |= MCIS_DISABLE_TX_VID_CHECK;
2608 } else {
2609     /*
2610      * Assert that the specified flags are consistent with the
2611      * flags specified by previous calls to mac_unicast_add().
2612      */
2613     ASSERT(((flags & MAC_UNICAST_TAG_DISABLE) != 0 &&
2614            (mcip->mci_state_flags & MCIS_TAG_DISABLE) != 0) ||
2615            ((flags & MAC_UNICAST_TAG_DISABLE) == 0 &&
2616            (mcip->mci_state_flags & MCIS_TAG_DISABLE) == 0));
2617
2618     ASSERT(((flags & MAC_UNICAST_STRIP_DISABLE) != 0 &&
2619            (mcip->mci_state_flags & MCIS_STRIP_DISABLE) != 0) ||
2620            ((flags & MAC_UNICAST_STRIP_DISABLE) == 0 &&
2621            (mcip->mci_state_flags & MCIS_STRIP_DISABLE) == 0));
2622
2623     ASSERT(((flags & MAC_UNICAST_DISABLE_TX_VID_CHECK) != 0 &&
2624            (mcip->mci_state_flags & MCIS_DISABLE_TX_VID_CHECK) != 0) ||
2625            ((flags & MAC_UNICAST_DISABLE_TX_VID_CHECK) == 0 &&
2626            (mcip->mci_state_flags & MCIS_DISABLE_TX_VID_CHECK) == 0));
2627
2628     /*
2629      * Make sure the client is consistent about its requests
2630      * for MAC addresses. I.e. all requests from the clients
2631      * must have the MAC_UNICAST_HW flag set or clear.
2632      */
2633     if ((mcip->mci_state_flags & MCIS_UNICAST_HW) != 0 &&
2634         !is_unicast_hw ||
2635         (mcip->mci_state_flags & MCIS_UNICAST_HW) == 0 &&
2636         is_unicast_hw) {
2637         err = EINVAL;
2638         goto bail_out;
2639     }
2640 }
2641
2642     /*
2643      * Make sure the MAC address is not already used by
2644      * another MAC client defined on top of the same
2645      * underlying NIC. Unless we have MAC_CLIENT_FLAGS_MULTI_PRIMARY
2646      * set when we allow a passive client to be present which will
2647      * be activated when the currently active client goes away - this
2648      * works only with primary addresses.
2649     */
2650     if ((check_dups || is_primary || is_vnic_primary) &&
2651         mac_addr_in_use(mip, mac_addr, vid)) {
2652         /*
2653          * Must have set the multiple primary address flag when
2654          * we did a mac_client_open AND this should be a primary
2655          * MAC client AND there should not already be a passive
2656          * primary. If all is true then we let this succeed
2657          * even if the address is a dup.
2658        */
2659     if ((mcip->mci_flags & MAC_CLIENT_FLAGS_MULTI_PRIMARY) == 0 ||
2660         (mcip->mci_flags & MAC_CLIENT_FLAGS_PRIMARY) == 0 ||
2661         mac_get_passive_primary_client(mip) != NULL) {
2662         *diag = MAC_DIAG_MACADDR_INUSE;
2663         err = EEXIST;
2664         goto bail_out;
2665     }
2666     ASSERT((mcip->mci_flags &

```

```

2666             MAC_CLIENT_FLAGS_PASSIVE_PRIMARY) == 0);
2667     mcip->mci_flags |= MAC_CLIENT_FLAGS_PASSIVE_PRIMARY;
2668     kmem_free(mrp, sizeof (*mrp));
2669
2670     /*
2671      * Stash the unicast address handle, we will use it when
2672      * we set up the passive client.
2673      */
2674     mcip->mci_p_unicast_list = muip;
2675     *mah = (mac_unicast_handle_t)muip;
2676     return (0);
2677 }
2678
2679 err = mac_client_datapath_setup(mcip, vid, mac_addr, mrp,
2680     is_primary || is_vnic_primary, muip);
2681 if (err != 0)
2682     goto bail_out;
2683
2684 kmem_free(mrp, sizeof (*mrp));
2685 *mah = (mac_unicast_handle_t)muip;
2686 return (0);
2687
2688 bail_out:
2689     if (fastpath_disabled)
2690         mac_fastpath_enable((mac_handle_t)mip);
2691     if (mcip->mci_state_flags & MCIS_EXCLUSIVE) {
2692         mip->mi_state_flags &= ~MIS_EXCLUSIVE;
2693         if (mip->mi_state_flags & MIS_LEGACY) {
2694             mip->mi_capab_legacy.ml_active_clear(
2695                 mip->mi_driver);
2696         }
2697     }
2698     kmem_free(mrp, sizeof (*mrp));
2699     kmem_free(muip, sizeof (mac_unicast_impl_t));
2700 }
2701 }
2702
2703 unchanged_portion_omitted

```

```
*****
19026 Thu Feb 20 18:59:07 2014
new/usr/src/uts/common/sys/mac.h
2553 mac address should be a dladm link property
*****
_____ unchanged_portion_omitted_


136 typedef struct mac_addrprop_s {
137     uint32_t          ma_len;
138     uint8_t           ma_addr[MAXMACADDRLEN];
139 } mac_addrprop_t;

141 #endif /* ! codereview */
142 #define MAXLINKPROPNAME      256           /* max property name len */

144 /*
145 * Public properties.
146 *
147 * Note that there are 2 sets of parameters: the *_EN_* values are
148 * those that the Administrator configures for autonegotiation. The
149 * _ADV_* values are those that are currently exposed over the wire.
150 */
151 typedef enum {
152     MAC_PROP_DUPLEX = 0x00000001,
153     MAC_PROP_SPEED,
154     MAC_PROP_STATUS,
155     MAC_PROP_AUTONEG,
156     MAC_PROP_EN_AUTONEG,
157     MAC_PROP_MTU,
158     MAC_PROP_ZONE,
159     MAC_PROP_AUTOPUSH,
160     MAC_PROP_FLOWCTRL,
161     MAC_PROP_ADV_1000FDX_CAP,
162     MAC_PROP_EN_1000FDX_CAP,
163     MAC_PROP_ADV_1000HDX_CAP,
164     MAC_PROP_EN_1000HDX_CAP,
165     MAC_PROP_ADV_100FDX_CAP,
166     MAC_PROP_EN_100FDX_CAP,
167     MAC_PROP_ADV_100HDX_CAP,
168     MAC_PROP_EN_100HDX_CAP,
169     MAC_PROP_ADV_10FDX_CAP,
170     MAC_PROP_EN_10FDX_CAP,
171     MAC_PROP_ADV_10HDX_CAP,
172     MAC_PROP_EN_10HDX_CAP,
173     MAC_PROP_ADV_100T4_CAP,
174     MAC_PROP_EN_100T4_CAP,
175     MAC_PROP_IPTUN_HOLIMIT,
176     MAC_PROP_IPTUN_ENCAPLIMIT,
177     MAC_PROP_WL_ESSID,
178     MAC_PROP_WL_BSSID,
179     MAC_PROP_WL_BSSTYPE,
180     MAC_PROP_WL_LINKSTATUS,
181     MAC_PROP_WL_DESIRED_RATES,
182     MAC_PROP_WL_SUPPORTED_RATES,
183     MAC_PROP_WL_AUTH_MODE,
184     MAC_PROP_WL_ENCRYPTION,
185     MAC_PROP_WL_RSSI,
186     MAC_PROP_WL_PHY_CONFIG,
187     MAC_PROP_WL_CAPABILITY,
188     MAC_PROP_WL_WPA,
189     MAC_PROP_WL_SCANRESULTS,
190     MAC_PROP_WL_POWER_MODE,
191     MAC_PROP_WL_RADIO,
192     MAC_PROP_WL_ESS_LIST,
193     MAC_PROP_WL_KEY_TAB,
194     MAC_PROP_WL_CREATE_IBSS,
```

```
195     MAC_PROP_WL_SETOPTIE,
196     MAC_PROP_WL_DELKEY,
197     MAC_PROP_WL_KEY,
198     MAC_PROP_WL_MLME,
199     MAC_PROP_TAGMODE,
200     MAC_PROP_ADV_10GFDX_CAP,
201     MAC_PROP_EN_10GFDX_CAP,
202     MAC_PROP_PVID,
203     MAC_PROP_LLIMIT,
204     MAC_PROP_LDECAY,
205     MAC_PROP_RESOURCE,
206     MAC_PROP_RESOURCE_EFF,
207     MAC_PROP_RXRINGS RANGE,
208     MAC_PROP_TXRINGS RANGE,
209     MAC_PROP_MAX_TX_RINGS_AVAIL,
210     MAC_PROP_MAX_RX_RINGS_AVAIL,
211     MAC_PROP_MAX_RXHWCLNT_AVAIL,
212     MAC_PROP_MAX_TXHWCLNT_AVAIL,
213     MAC_PROP_IB_LINKMODE,
214     MAC_PROP_MACADDRESS,
215 #endif /* ! codereview */
216     MAC_PROP_PRIVATE = -1
217 } mac_prop_id_t;

219 /*
220 * Flags to figure out r/w status of legacy ndd props.
221 */
222 #define MAC_PROP_PERM_READ          0x0001
223 #define MAC_PROP_PERM_WRITE         0x0010
224 #define MAC_PROP_MAP_KSTAT          0x0100
225 #define MAC_PROP_PERM_RW            (MAC_PROP_PERM_READ|MAC_PROP_PERM_WRITE)
226 #define MAC_PROP_FLAGS_RK           (MAC_PROP_PERM_READ|MAC_PROP_MAP_KSTAT)

228 #ifdef __KERNEL__

230 /*
231 * There are three ranges of statistics values. 0 to 1 - MAC_STAT_MIN are
232 * interface statistics maintained by the mac module. MAC_STAT_MIN to 1 -
233 * MACTYPE_STAT_MIN are common MAC statistics defined by the mac module and
234 * maintained by each driver. MACTYPE_STAT_MIN and above are statistics
235 * defined by MAC-Type plugins and maintained by each driver.
236 */
237 #define MAC_STAT_MIN               1000
238 #define MACTYPE_STAT_MIN           2000

240 #define IS_MAC_STAT(stat)          \
241     (stat >= MAC_STAT_MIN && stat < MACTYPE_STAT_MIN)
242 #define IS_MACTYPE_STAT(stat)      (stat >= MACTYPE_STAT_MIN)

244 /*
245 * Statistics maintained by the mac module, and possibly populated as link
246 * statistics.
247 */
248 enum mac_mod_stat {
249     MAC_STAT_LINK_STATE,
250     MAC_STAT_LINK_UP,
251     MAC_STAT_PROMISC,
252     MAC_STAT_LOWLINK_STATE,
253     MAC_STAT_HDROPS
254 };

256 /*
257 * Do not reorder, and add only to the end of this list.
258 */
259 enum mac_driver_stat {
260     /* MIB-II stats (RFC 1213 and RFC 1573) */
```

```

261     MAC_STAT_IFSPEED = MAC_STAT_MIN,
262     MAC_STAT_MULTIRCV,
263     MAC_STAT_BRDCSTRCV,
264     MAC_STAT_MULTIXMT,
265     MAC_STAT_BRDCSTXMT,
266     MAC_STAT_NORCVBUF,
267     MAC_STAT_IERRORS,
268     MAC_STAT_UNKNOWNS,
269     MAC_STAT_NOXMTBUF,
270     MAC_STAT_OERRORS,
271     MAC_STAT_COLLISIONS,
272     MAC_STAT_RBYTES,
273     MAC_STAT_IPACKETS,
274     MAC_STAT_OBYTES,
275     MAC_STAT_OPACKETS,
276     MAC_STAT_UNDERFLOWS,
277     MAC_STAT_OVERFLOWS
278 };
280 #define MAC_NSTAT      (MAC_STAT_OVERFLOWS - MAC_STAT_IFSPEED + 1)

282 #define MAC_STAT_ISACOUNTER(_stat) \
283     (_stat) == MAC_STAT_MULTIRCV || \
284     (_stat) == MAC_STAT_BRDCSTRCV || \
285     (_stat) == MAC_STAT_MULTIXMT || \
286     (_stat) == MAC_STAT_BRDCSTXMT || \
287     (_stat) == MAC_STAT_NORCVBUF || \
288     (_stat) == MAC_STAT_IERRORS || \
289     (_stat) == MAC_STAT_UNKNOWNS || \
290     (_stat) == MAC_STAT_NOXMTBUF || \
291     (_stat) == MAC_STAT_OERRORS || \
292     (_stat) == MAC_STAT_COLLISIONS || \
293     (_stat) == MAC_STAT_RBYTES || \
294     (_stat) == MAC_STAT_IPACKETS || \
295     (_stat) == MAC_STAT_OBYTES || \
296     (_stat) == MAC_STAT_OPACKETS || \
297     (_stat) == MAC_STAT_UNDERFLOWS || \
298     (_stat) == MAC_STAT_OVERFLOWS)

300 /*
301 * Immutable information. (This may not be modified after registration).
302 */
303 typedef struct mac_info_s {
304     uint_t          mi_media;
305     uint_t          mi_nativemedia;
306     uint_t          mi_addr_length;
307     uint8_t         *mi_unicst_addr;
308     uint8_t         *mi_brdcst_addr;
309 } mac_info_t;

311 /*
312 * When VNICs are created on top of the NIC, there are two levels
313 * of MAC layer, a lower MAC, which is the MAC layer at the level of the
314 * physical NIC, and an upper MAC, which is the MAC layer at the level
315 * of the VNIC. Each VNIC maps to a MAC client at the lower MAC, and
316 * the SRS and classification is done at the lower MAC level. The upper
317 * MAC is therefore for the most part pass-through, and therefore
318 * special processing needs to be done at the upper MAC layer when
319 * dealing with a VNIC.
320 */
321 /* This capability allows the MAC layer to detect when a VNIC is being
322 * accessed, and implement the required shortcuts.
323 */
325 typedef void *(*mac_client_handle_fn_t)(void *);
```

```

327 typedef struct mac_capab_vnic_s {
328     void             *mcv_arg;
329     mac_client_handle_fn_t  mcv_mac_client_handle;
330 } mac_capab_vnic_t;

332 typedef void (*mac_rename_fn_t)(const char *, void *);
333 typedef mblk_t *(*mac_tx_ring_fn_t)(void *, mblk_t *, uintptr_t,
334                                     mac_ring_handle_t *);
335 typedef struct mac_capab_aggr_s {
336     mac_rename_fn_t mca_rename_fn;
337     int (*mca_unicst)(void *, const uint8_t *);
338     mac_tx_ring_fn_t mca_find_tx_ring_fn;
339     void *mca_arg;
340 } mac_capab_aggr_t;

342 /* Bridge transmit and receive function signatures */
343 typedef mblk_t *(*mac_bridge_tx_t)(mac_handle_t, mac_ring_handle_t, mblk_t *);
344 typedef void (*mac_bridge_rx_t)(mac_handle_t, mac_resource_handle_t, mblk_t *);
345 typedef void (*mac_bridge_ref_t)(mac_handle_t, boolean_t);
346 typedef link_state_t (*mac_bridge_ls_t)(mac_handle_t, link_state_t);

348 /* must change mac_notify_cb_list[] in mac_provider.c if this is changed */
349 typedef enum {
350     MAC_NOTE_LINK,
351     MAC_NOTE_UNICST,
352     MAC_NOTE_TX,
353     MAC_NOTE_DEVPROMISC,
354     MAC_NOTE_FASTPATH_FLUSH,
355     MAC_NOTE_SDU_SIZE,
356     MAC_NOTE_DEST,
357     MAC_NOTE_MARGIN,
358     MAC_NOTE_CAPAB_CHG,
359     MAC_NOTE_LOWLINK,
360     MAC_NOTE_ALLOWED_IPS,
361     MAC_NNOTE /* must be the last entry */
362 } mac_notify_type_t;

364 typedef void      (*mac_notify_t)(void *, mac_notify_type_t);
365 typedef void      (*mac_rx_t)(void *, mac_resource_handle_t, mblk_t *,
366                             boolean_t);
367 typedef mblk_t    *(*mac_receive_t)(void *, int);

369 /*
370 * MAC resource types
371 */
372 typedef enum {
373     MAC_RX_FIFO = 1
374 } mac_resource_type_t;

376 typedef int      (*mac_intr_enable_t)(mac_intr_handle_t);
377 typedef int      (*mac_intr_disable_t)(mac_intr_handle_t);

379 typedef struct mac_intr_s {
380     mac_intr_handle_t    mi_handle;
381     mac_intr_enable_t   mi_enable;
382     mac_intr_disable_t  mi_disable;
383     ddi_intr_handle_t   mi_ddi_handle;
384     boolean_t            mi_ddi_shared;
385 } mac_intr_t;

387 typedef struct mac_rx_fifo_s {
388     mac_resource_type_t  mrf_type; /* MAC_RX_FIFO */
389     mac_intr_t           mrf_intr;
390     mac_receive_t        mrf_receive;
391     void                *mrf_rx_arg;
392     uint32_t             mrf_flow_priority;
```

```

393     /*
394      * The CPU this flow is to be processed on. With intrd and future
395      * things, we should know which CPU the flow needs to be processed
396      * and get a squeue assigned on that CPU.
397      */
398     uint_t          mrf_cpu_id;
399 } mac_rx_fifo_t;

401 #define mrf_intr_handle    mrf_intr.mi_handle
402 #define mrf_intr_enable     mrf_intr.mi_enable
403 #define mrf_intr_disable    mrf_intr.mi_disable

405 typedef union mac_resource_u {
406     mac_resource_type_t   mr_type;
407     mac_rx_fifo_t        mr_fifo;
408 } mac_resource_t;

410 typedef enum {
411     MAC_ADDRTYPE_UNICAST,
412     MAC_ADDRTYPE_MULTICAST,
413     MAC_ADDRTYPE_BROADCAST
414 } mac_addrtype_t;

416 typedef struct mac_header_info_s {
417     size_t           mhi_hdrlen;
418     size_t           mhi_pktsize;
419     const uint8_t   *mhi_daddr;
420     const uint8_t   *mhi_saddr;
421     uint32_t         mhi_origsap;
422     uint32_t         mhi_bindsap;
423     mac_addrtype_t mhi_dsttype;
424     uint16_t         mhi_tci;
425     boolean_t        mhi_istagged;
426     boolean_t        mhi_ispvrid;
427 } mac_header_info_t;

429 /*
430  * Function pointer to match dls client signature. Should be same as
431  * dls_rx_t to allow a soft ring to bypass DLS layer and call a DLS
432  * client directly.
433 */
434 typedef void        (*mac_direct_rx_t)(void *, mac_resource_handle_t,
435                                         mblk_t *, mac_header_info_t);

437 typedef mac_resource_handle_t  (*mac_resource_add_t)(void *, mac_resource_t *);
438 typedef int          (*mac_resource_bind_t)(void *,
439                                              mac_resource_handle_t, processorid_t);
440 typedef void        (*mac_resource_remove_t)(void *, void *);
441 typedef void        (*mac_resource_quiesce_t)(void *, void *);
442 typedef void        (*mac_resource_restart_t)(void *, void *);
443 typedef int          (*mac_resource_modify_t)(void *, void *,
444                                              mac_resource_t *);
445 typedef void        (*mac_change_upcall_t)(void *, mac_direct_rx_t,
446                                         void *);

448 /*
449  * MAC-Type plugin interfaces
450 */
451
452 typedef int          (*mtops_addr_verify_t)(const void *, void *);
453 typedef boolean_t    (*mtops_sap_verify_t)(uint32_t, uint32_t *, void *);
454 typedef mblk_t       (*(*mtops_header_t)(const void *, const void *,
455                                         uint32_t, void *, mblk_t *, size_t));
456 typedef int          (*mtops_header_info_t)(mblk_t *, void *,
457                                              mac_header_info_t *);
458 typedef boolean_t    (*mtops_pdata_verify_t)(void *, size_t);

```

```

459 typedef mblk_t          (*(*mtops_header_modify_t)(mblk_t *, void *));
460 typedef void             (*mtops_link_details_t)(char *, size_t, mac_handle_t,
461                                               void *);
462
463 typedef struct mactype_ops_s {
464     uint_t               mtops_ops;
465     /*
466      * mtops_unicst_verify() returns 0 if the given address is a valid
467      * unicast address, or a non-zero errno otherwise.
468      */
469     mtops_addr_verify_t  mtops_unicst_verify;
470     /*
471      * mtops_multicst_verify() returns 0 if the given address is a
472      * valid multicast address, or a non-zero errno otherwise. If the
473      * media doesn't support multicast, ENOTSUP should be returned (for
474      * example).
475      */
476     mtops_addr_verify_t  mtops_multicst_verify;
477     /*
478      * mtops_sap_verify() returns B_TRUE if the given SAP is a valid
479      * SAP value, or B_FALSE otherwise.
480      */
481     mtops_sap_verify_t   mtops_sap_verify;
482     /*
483      * mtops_header() is used to allocate and construct a MAC header.
484      */
485     mtops_header_t        mtops_header;
486     /*
487      * mtops_header_info() is used to gather information on a given MAC
488      * header.
489      */
490     mtops_header_info_t   mtops_header_info;
491     /*
492      * mtops_pdata_verify() is used to verify the validity of MAC
493      * plugin data. It is called by mac_register() if the driver has
494      * supplied MAC plugin data, and also by mac_pdata_update() when
495      * drivers update the data.
496      */
497     mtops_pdata_verify_t  mtops_pdata_verify;
498     /*
499      * mtops_header_cook() is an optional callback that converts (or
500      * "cooks") the given raw header (as sent by a raw DLPI consumer)
501      * into one that is appropriate to send down to the MAC driver.
502      * Following the example above, an Ethernet header sent down by a
503      * DLPI consumer would be converted to whatever header the MAC
504      * driver expects.
505      */
506     mtops_header_modify_t mtops_header_cook;
507     /*
508      * mtops_header_uncook() is an optional callback that does the
509      * opposite of mtops_header_cook(). It "uncooks" a given MAC
510      * header (as received from the driver) for consumption by raw DLPI
511      * consumers. For example, for a non-Ethernet plugin that wants
512      * raw DLPI consumers to be fooled into thinking that the device
513      * provides Ethernet access, this callback would modify the given
514      * mblk_t such that the MAC header is converted to an Ethernet
515      * header.
516      */
517     mtops_header_modify_t mtops_header_uncook;
518     /*
519      * mtops_link_details() is an optional callback that provides
520      * extended information about the link state. Its primary purpose
521      * is to provide type-specific support for syslog contents on
522      * link up events. If no implementation is provided, then a default
523      * implementation will be used.
524 */

```

```

525     mtops_link_details_t    mtops_link_details;
526 } mactype_ops_t;

528 /*
529  * mtops_ops exists for the plugin to enumerate the optional callback
530  * entrypoints it has defined. This allows the mac module to define
531  * additional plugin entrypoints in mactype_ops_t without breaking backward
532  * compatibility with old plugins.
533 */
534 #define MTOPS_PDATA_VERIFY      0x001
535 #define MTOPS_HEADER_COOK       0x002
536 #define MTOPS_HEADER_UNCOOK     0x004
537 #define MTOPS_LINK_DETAILS      0x008

539 /*
540  * Provide mapping for legacy ndd ioctls relevant to that mactype.
541  * Note that the ndd ioctls are obsolete, and may be removed in a future
542  * release of Solaris. The ndd ioctls are not typically used in legacy
543  * ethernet drivers. New datalink drivers of all link-types should use
544  * dladm(lm) interfaces for administering tunables and not have to provide
545  * a mapping.
546 */
547 typedef struct mac_ndd_mapping_s {
548     char          *mp_name;
549     union {
550         mac_prop_id_t   u_id;
551         uint_t          u_kstat;
552     } u_mp_id;
553     long           mp_minval;
554     long           mp_maxval;
555     size_t          mp_valsize;
556     int            mp_flags;
557 } mac_ndd_mapping_t;

558 #define mp_prop_id      u_mp_id.u_id
559 #define mp_kstat        u_mp_id.u_kstat

562 typedef struct mac_stat_info_s {
563     uint_t          msi_stat;
564     char           *msi_name;
565     uint_t          msi_type;      /* as defined in kstat_named_init(9F) */
566     uint64_t        msi_default;
567 } mac_stat_info_t;

569 typedef struct mactype_register_s {
570     uint_t          mtr_version;    /* set by mactype_alloc() */
571     const char     *mtr_ident;
572     mactype_ops_t  *mtr_ops;
573     uint_t          mtr_mactype;
574     uint_t          mtr_nativetype;
575     uint_t          mtr_addrlen;
576     uint8_t         *mtr_brdcst_addr;
577     mac_stat_info_t *mtr_stats;
578     size_t          mtr_statcount;
579     mac_ndd_mapping_t *mtr_mapping;
580     size_t          mtr_mappingcount;
581 } mactype_register_t;

583 /*
584  * Driver interface functions.
585 */
586 extern int           mac_open_by_linkid(datalink_id_t,
587                                         mac_handle_t *);
588 extern int           mac_open_by_linkname(const char *,
589                                         mac_handle_t *);
590 extern const char    *mac_name(mac_handle_t);

```

```
657 }  
658 #endif  
660 #endif /* _SYS_MAC_H */
```