



Introduction to Microsemi Ethernet Switch API

User Guide

Table of Contents

1. Introduction	3
1.1. Scope.....	3
1.2. Audience.....	3
1.3. Prerequisites	4
2. Architecture	4
2.1. The big picture.....	4
2.1.1. Main Application	4
2.1.2. MEBA	5
2.1.3. MESA	5
2.1.4. Third-party Application	5
2.1.5. Linux Kernel	5
3. MESA source introduction.....	7
3.1. Co-Existing with Unified-API.....	7
3.2. Capabilities	7
3.2.1. Availability of functions.....	8
3.2.2. Availability of <code>struct</code> members	10
3.2.3. Fixed length arrays	13
3.2.4. Case study - QOS.....	16
4. Build system, auto-generation and migration	19
4.1. Development environment	19
4.2. Building the sources	20
4.2.1. Auto generated content.....	21
5. Application migration	22
6. Reference boards and demo system	23
6.1. CLI Application.....	23
6.2. Building the demo application	23
6.3. Download.....	23
6.4. Running	24
6.5. Adding a Test	26
Appendix A: C++ CapArray	27
References.....	37

1. Introduction

MSCC-ENT offers a single API to access the functionality of ethernet switches and PHY's. Currently this API offers two different C interfaces, the Unified-API (which has existed for a long time), and the MESA layer. Both layers are equally supported, and customer can choose to use the layer which fits best into their design.

The difference between the Unified-API and MESA, is that public header files part of the Unified-API is heavily annotated with pre-processor statements (`#ifdef`, `#else` and `#endif`). These pre-processor statements is controlled by compiler defines, that is specifying a single chip (along with a list of features) supported by the binary library. These pre-processor statements is causing both the application programming interface (API) and application binary interface (ABI) of the library to differ for each supported chip and the optional features.

The MESA interface on the other hand is offering a stable API/ABI across all targets (and options) supported by a given release. Instead of using pre-processor statements the MESA interface is using capabilities to annotate what features is available in a given instance of the library. The MESA interface is thereby allowing customers to only build a single application, and link (load- or compile-time) it with the MESA library supporting a given chip.

Currently the MESA interface, is implemented as a thin layer on-top of the Unified-API library. The MESA layer is only converting between the stable API/ABI offered by the MESA headers, to the chip specific API/ABI offered by a given compilation of the Unified-API.

This means that the Unified-API and the MESA interfaces is offering the same level of abstraction. Documentation written for the Unified-API will also be useful if using the MESA API.

1.1. Scope

This document is intended as an introduction document of the MESA interface, not a reference or description of all the API's included as part of the MESA interface.

The document will describe the general terms/facilities that is available when the MESA interface is used.

Beside, from the core MESA library, the document also describe the `mesa_demo` application, which can be used to experiment with the API, using one of the reference boards.

1.2. Audience

The intended audience for this document is software developers who need to understand the facilities offered by the MESA interface. The document is expected to be useful in the following scenario's:

- Customers who is new to the API SW offered by MSCC-ENT, and wants to evaluate it for specific product/purpose.

- Customers who is starting up a new project based on the MESA interface.
- Customers who is considering changing from the Unified-API to the MESA interface.

1.3. Prerequisites

This document assumes the reader possesses the following skills/resources:

1. Fluent in C and to some extent Makefiles.
2. Root access to a recent Linux development environment, and experienced in working with a Linux shell.
 - a. Building new images from the sources (including boot-loader, BSP and application) requires a 64-bit Linux machine with at least 8GB of RAM, 50GB of disk space and 4 CPU cores. This document uses Ubuntu 16.04LTS as reference.
 - b. Access to a TFTP server that can be used for SW upgrades.
 - c. RS232 terminal to access the target.
3. MSCC API source package.
4. MSCC reference board (with WebStaX/SMBStaX/IStaX/CEServices 4.0 or newer).

2. Architecture

The MESA library is a driver which can help the application configuring the switch core and/or read status. The driver contains no threads, it implements no protocols and have no awareness of network frames. Without an application, the driver can not do anything.

The driver is OS agnostic, and requires that the application provide function pointers for accessing registers, locking, tracing etc.

2.1. The big picture

The MESA API is just a single component in a much larger system. This section will try to illustrate how/where the API fits in to the system architecture. This section will assume that the target system is running Linux.

2.1.1. Main Application

The application in the system which is linking in the MEBA/MESA libraries is referred to as the main application. This application is exclusively owning the API, and can use it to build the needed functionality. The main application will often expose a number of machine and/or human targeted interfaces that the outside world can use to apply configuration or query status. In turnkey systems delivered by MSCC, the WebStaX application is typically the main application, while in API projects, the customer either have an existing application or will design an application which owns the API. The MESA project does include a small and simple sample main-application which shows how to instantiate the API, and to setup simple stuff, this is the mini-application. The mini application is covered in more details later. See [ug1068] for more details.

2.1.2. MEBA

The MSCC Ethernet Board API (MEBA) is created to provide an abstract interface to all the board facilities. MSCC provides a MEBA library for each of the reference systems, and the main application should use it to handle the board specific details. If customers are doing a custom board, then they should also do an implementation of the MEBA library which matches their board. See [ug1069] for more details.

2.1.3. MESA

The MSCC Ethernet Switch API (MESA) is the component covered in this document.

2.1.4. Third-party Application

A system typically contains many applications running on the same CPU (either the integrated CPU or an external CPU), this is well supported. But there is only a single application which can instantiate the API, the main application. All other application is being referred to as third-party applications. Third-party applications can communicate with the Linux kernel, with other application, but if they need resources from MESA/MEBA, then it must go through the main application. In the WebStaX system, the WebStaX application is offering a JSON-IPC connection to third-party application which wants to add certain functionality.

2.1.5. Linux Kernel

The Linux kernel which can be found in the BSP, is offering kernel drivers for a number of the CPU peripherals (assuming the integrated CPU is used). The following subsections is elaborating a bit more on these facilities. See [ug1068] for more details.

2.1.5.1. I2C/SPI/UIO

When the MEBA/MESA libraries are instantiated, then function pointers to access register in the various devices must be provided. This is how these libraries can be OS agnostic.

To reach the register in the various peripherals, the switch-core is offering I2C and SPI controllers which is connected to the peripherals by the board. To reach the registers in the SwitchCore it self a UIO driver is exposed.

These buses/devices is typically exposed as character devices in the `/dev/` file system by the kernel. They can either be memory mapped, or the application can issues read/write commands to control the bus.

The main application can then use these devices, and provide function pointers for reading/writing the needed registers.

2.1.5.2. NIC(s)

The Switch Core have the concept of a CPU port, which allows frames to be moved between the CPU and the Switch core. Frames going in/out if the CPU port (typically) have an interface-frame-header which carry information on why the packet is being copied/moved to the CPU, or how they should be injected in the Switch Core.

The MSCC BSP will expose the CPU port as a normal NIC interface. The NIC driver will expose the frames as-is without doing any processing. This means that they will include the IFH, which carry important information. The application running in user-space can then open a raw socket and receive the frames, process them, (set up the MESA API if needed), and inject frames as a response.

2.1.5.3. NAND/NOR Drivers

The kernel also provides NAND/NOR drivers, which is exposed as block devices. A file system is typically created on-top of the block devices, which then can be used from user-space applications.

3. MESA source introduction

The key difference between the Unified-API and the MESA library, is that the MESA library is API/ABI identical for a given release. This means that a single application can be build once, and linked with the API matching the target.

Most of the API definitions in Unified-API and MESA, are very similar. In most cases the differences follows the rules explained in this section.

3.1. Co-Existing with Unified-API

The MESA library and the Unified-API are mutually exclusive from an application point of view. The application must choose if it want to use the one or the other, and only include header files from the desired layer.

3.2. Capabilities

The Unified-API is heavily annotated with `ifdef` which specifies what facilities is available in a given configuration. The MESA library does not use these pre-processor statements as they generally breaks API/ABI compatibility between different targets. Instead of using `ifdef`, the MESA layer uses capabilities.

A capability is an enumerated value, representing a given feature (see `mesa_cap_t`). Along with the list of defined capabilities, the MESA library implements a function to query the status of a given capability:

```
uint32_t mesa_capability(mesa_inst_t inst, int cap);
```

mscc/ethernet/switch/api/capability.h

```

1  typedef enum {
2      MESA_CAP_MISC_GPIO_CNT = 0,    /**< Number of GPIOs */
3      MESA_CAP_MISC_SGPIO_CNT,     /**< Number of SGPIO groups */
4      MESA_CAP_MISC_PORT_GPIO,     /**< Port GPIO */
5      MESA_CAP_MISC_INTERRUPTS,    /**< Interrupts */
6      MESA_CAP_MISC_IRQ_CONTROL,   /**< IRQ control */
7      MESA_CAP_MISC_FAN,           /**< Fan control */
8
9      // Port
10     MESA_CAP_PORT_CNT = 100,      /**< Maximum number of ports */
11
12     // Many more
13 } mesa_cap_t;
14
15 uint32_t mesa_capability(mesa_inst_t inst, int cap);

```



The capabilities is used to signals booleans (if something is there or not) or to specify number (how many instances does the chip support).

The following part of this section explains how the capability system is used to achieve the stable API/ABI.

3.2.1. Availability of functions

There are many examples of complete function groups, which only exists in certain targets. Either it is because new features are added in the chips, or because a given component is completely redesigned in such way that it is not possible to backwards compatible.

An example of this can be found in `vtss_afi_api.h`. A small simplified snippet of that file which illustrates how the pre-processors is used currently is shown below:

vtss_afi_api.h

```
1 typedef u32 vtss_afi_id_t;
2 #if defined(VTSS_AFI_V1)
3 typedef struct {
4     u32 fps;
5     u32 actual_fps;
6 } vtss_afi_frm_dscr_t;
7 vtss_rc vtss_afi_alloc(const vtss_inst_t inst, vtss_afi_frm_dscr_t *const dscr,
8                       vtss_afi_id_t *const id);
9 vtss_rc vtss_afi_free(const vtss_inst_t inst, vtss_afi_id_t id);
10 #endif /* defined(VTSS_AFI_V1) */
11
12 #if defined(VTSS_AFI_V2)
13 typedef u32 vtss_afi_fastid_t;
14
15 typedef struct {
16     vtss_port_no_t port_no;
17     vtss_prio_t prio;
18     u32 frm_cnt;
19 } vtss_afi_fast_inj_alloc_cfg_t;
20
21 vtss_rc vtss_afi_fast_inj_alloc(const vtss_inst_t inst,
22                                const vtss_afi_fast_inj_alloc_cfg_t *const cfg,
23                                vtss_afi_fastid_t *const fastid);
24
25 vtss_rc vtss_afi_fast_inj_free(const vtss_inst_t inst, vtss_afi_fastid_t fastid);
26 #endif // VTSS_AFI_V2
```


vtss_afi_api.h

```

1  typedef uint32_t mesa_afi_id_t CAP(AFI_V1);
2  typedef struct {
3      uint32_t fps;
4  } mesa_afi_frm_dscr_t CAP(AFI_V1);
5  typedef struct {
6      uint32_t fps;
7      u32 actual_fps;
8  } mesa_afi_frm_dscr_actual_t CAP(AFI_V1);
9
10 mesa_rc mesa_afi_alloc(const mesa_inst_t      inst,
11                       const mesa_afi_frm_dscr_t *const dscr,
12                       mesa_afi_frm_dscr_actual_t *const actual,
13                       mesa_afi_id_t           *const id) CAP(AFI_V1);
14
15 mesa_rc mesa_afi_free(const mesa_inst_t inst,
16                      mesa_afi_id_t     id) CAP(AFI_V1);
17
18 typedef uint32_t mesa_afi_fastid_t CAP(AFI_V2);
19 typedef struct {
20     mesa_port_no_t port_no;
21     mesa_prio_t prio;
22     uint32_t frm_cnt;
23 } mesa_afi_fast_inj_alloc_cfg_t CAP(AFI_V2);
24
25 mesa_rc mesa_afi_fast_inj_alloc(
26     const mesa_inst_t      inst,
27     const mesa_afi_fast_inj_alloc_cfg_t *const cfg,
28     mesa_afi_fastid_t      *const fastid) CAP(AFI_V2);
29
30 mesa_rc mesa_afi_fast_inj_free(const mesa_inst_t      inst,
31                               mesa_afi_fastid_t fastid) CAP(AFI_V2);

```

Instead of having either `VTSS_AFI_V1` or `VTSS_AFI_V2` the header files now contains the prototypes of both. The source code is annotated with `CAP(XXX)` which is a macro expanding to nothing. The purpose of this is to document what capability the application should use to check if a given group of functions are supported or not.

An example of this is show here:

Application code with ifdef's

```

1  #if defined(VTSS_AFI_V1)
2      vtss_afi_alloc(inst, ...);
3  #elif defined(VTSS_AFI_V2)
4      vtss_afi_v2_fast_inj_free(inst, ...);
5  #endif

```

Application code without ifdef's

```

1  if (mesa_ent_api_capability(inst, MESA_CAP_AFI_V1)) {
2      mesa_afi_alloc(inst, ...);
3  } else if (mesa_ent_api_capability(inst, MESA_CAP_AFI_V2)) {
4      mesa_afi_fast_inj_alloc(inst, ...);
5  }

```

If a function is called, even though it is not implemented (according to the capability system), then it must return `MESA_RC_NOT_IMPLEMENTED`.

3.2.2. Availability of `struct` members

In the Unified API the pre-processor is often used to express small variations of the different targets supported by the API. This is an efficient way of expressing these differences, but the downside is that the application needs to have the same level of `ifdef` abstractions and thereby causing the API and APPL to be very closely bound together.

Following is an example of this pattern:

vtss_afi_api.h

```

1  typedef struct {
2      vtss_vcap_vr_t dscp;
3  #if defined(VTSS_FEATURE_ECE_KEY_IP_PROTO)
4      vtss_vcap_bit_t fragment;
5      vtss_vcap_u8_t proto;
6  #endif
7  #if defined(VTSS_FEATURE_ECE_KEY_SIP)
8      vtss_vcap_ip_t sip;
9  #endif
10 #if defined(VTSS_FEATURE_ECE_KEY_DIP)
11     vtss_vcap_ip_t dip;
12 #endif
13 #if defined(VTSS_FEATURE_ECE_KEY_SPORT)
14     vtss_vcap_vr_t sport;
15 #endif
16 #if defined(VTSS_FEATURE_ECE_KEY_DPORT)
17     vtss_vcap_vr_t dport;
18 #endif
19 } vtss_ece_frame_ipv4_t;

```

In reality, most targets actually supports most the keys. The `ifdef` will allow saving a few bytes, but a detailed analysis of this specific case showed that the `dip` field was the only difference when considering the current set of targets.

In the cases where we find that the actual difference in sizes, can be neglected, then we simply let the MESA header files include the super set of all members.

This is illustrated below:

mscc/ethernet/switch/api/evc.h

```
1 typedef struct {
2     mesa_vcap_vr_t  dscp;
3     mesa_vcap_bit_t fragment;
4     mesa_vcap_u8_t  proto;
5     mesa_vcap_ip_t  sip;
6     mesa_vcap_ip_t  dip CAP(EVC_ECE_DIP);
7     mesa_vcap_vr_t  sport;
8     mesa_vcap_vr_t  dport;
9 } mesa_ece_frame_ipv4_t CAP(EVC_ECE_CNT);
```



The `CAP(x)` macro is again used to annotate what capability is used to signal the availability of the different members.

If this approach (pulling in the super set of members) will cause a significant overhead, then the structure needs to be split into smaller structures. The majority of the structures was designed in such a way that all members could be included without causing significant overhead.

The few cases where a super set of all flags, is posing too big overhead, needed to be handled differently. This is covered in the next section.

3.2.2.1. Structure splitting

Just adding the super set of all attributes is not always the right thing to do, especially if that is causing large memory overhead, or if the extra attributes is very special/exotic and only useful in rare cases.

An example of this can be found in the MAC table, snippet is included below:

vtss_l2_api.current.h

```
1 typedef struct {
2     vtss_vid_mac_t vid_mac;
3     BOOL destination[VTSS_PORT_ARRAY_SIZE];
4     BOOL copy_to_cpu;
5     BOOL locked;
6     BOOL aged;
7     #if defined(VTSS_FEATURE_MAC_CPU_QUEUE)
8     vtss_packet_rx_queue_t cpu_queue;
9     #endif
10
11     #if defined(VTSS_FEATURE_VSTAX_V2)
12     struct {
13         BOOL enable;
14         BOOL remote_entry;
15         vtss_vstax_upsid_t upsid;
16         vtss_vstax_upspn_t upspn;
17     } vstax2;
18     #endif
19
20     #if defined(VTSS_FEATURE_SEAMLESS_REDUNDANCY)
21     struct {
22         BOOL enable;
23         vtss_sr_stream_id_t id;
24     } sr;
25     #endif
26 } vtss_mac_table_entry_t;
27
28 vtss_rc vtss_mac_table_add(const vtss_inst_t inst,
29                             const vtss_mac_table_entry_t *const entry);
30
31 vtss_rc vtss_mac_table_get(const vtss_inst_t inst,
32                             const vtss_vid_mac_t *const vid_mac,
33                             vtss_mac_table_entry_t *const entry);
```

The `destination` will be translated to a `mesa_port_mask_t` as illustrated earlier. The `cpu_queue` will just be included always with a `MCAP` attribute, but the `vstax2` and `sr` are quite special and does consume some space - we do not want to included them every representation of a mac-table in user space.

mscc/ethernet/switch/api/l2.suggested.h

```

1  typedef struct {
2      mesa_vid_mac_t vid_mac;
3      mesa_port_mask_t destination;
4      BOOL copy_to_cpu;
5      BOOL locked;
6      BOOL aged;
7      mesa_packet_rx_queue_t cpu_queue MCAP(MAC_CPU_QUEUE);
8  } mesa_mac_table_entry_t;
9
10 struct {
11     BOOL enable;
12     BOOL remote_entry;
13     mesa_vstax_upsid_t upsid;
14     mesa_vstax_upspn_t upspn;
15 } mesa_mac_table_entry_vstax2_t;
16
17 struct {
18     BOOL enable;
19     mesa_sr_stream_id_t id;
20 } mesa_mac_table_entry_sr_t;
21
22 mesa_rc mesa_mac_table_add(const mesa_inst_t inst,
23                           const mesa_mac_table_entry_t *const entry);
24
25 mesa_rc mesa_mac_table_get(const mesa_inst_t inst,
26                           const mesa_vid_mac_t *const vid_mac,
27                           mesa_mac_table_entry_t *const entry);
28
29 mesa_rc mesa_mac_table_adv_add(const mesa_inst_t inst,
30                               const mesa_mac_table_entry_t *const entry,
31                               const mesa_mac_table_entry_vstax2_t *const vstax,
32                               const mesa_mac_table_entry_sr_t *const sr);
33
34 mesa_rc mesa_mac_table_adv_get(const mesa_inst_t inst,
35                               const mesa_vid_mac_t *const vid_mac,
36                               mesa_mac_table_entry_t *const entry,
37                               mesa_mac_table_entry_vstax2_t *const vstax,
38                               mesa_mac_table_entry_sr_t *const sr);

```

The `vstax2` and `sr` structures has been removed from the `mesa_mac_table_entry_t`, and the `mesa_mac_table_entry_t` has thereby been unified for all target platforms. Users that needs to control the `vstax2` and/or `sr` fields will have to use the `mesa_mac_table_adv_add` method instead of the `mesa_mac_table_add` method.

3.2.3. Fixed length arrays

Defines known at compile time can be used to specify arrays, but if the length is a runtime parameter, which is the case with return values from `mesa_capability`, then the storage for the array needs to be allocated dynamic.

The Unified-API generally do not allocate any memory after it has been instantiated, and the same is true for the MESA library. It is always the responsibility of the application to allocate/handle-error/free any needed memory resources. The MESA library does not change that.

3.2.3.1. Port masks

Many data types uses a pre-processor define to specify the length of an array. That will again cause the ABI of the header to depend on the actual configuration, because most platforms has a different port count.

An example of this can be see below:

vtss_l2_api.h

```

1  typedef struct {
2  |   BOOL port_list[VTSS_PORT_ARRAY_SIZE];
3  |   vtss_vce_mac_t mac;
4  |   vtss_vce_tag_t tag;
5  |   vtss_vce_type_t type;
6  |   union { ... } frame;
7  | } vtss_vce_key_t;
```

Below is illustrated how the same behavioral can be achieved without making the ABI depending on the actual configuration.

mscc/ethernet/switch/api/l2.h

```

1  typedef struct {
2  |   mesa_port_mask_t port_list;
3  |   vtss_vce_mac_t mac;
4  |   vtss_vce_tag_t tag;
5  |   vtss_vce_type_t type;
6  |   union { ... } frame;
7  | } vtss_vce_key_t;
```

Instead of using an array of booleans, a port-mask is used instead. The port-mask can be dimensioned for worst case (current implementation is using 64bits, supporting up-to 64 ports), still be more space effective than the array of booleans.



An ABI compatible C header files can be used to emulate the old behavioral such that users in a C context will not notice that they are operating on a mask instead of an array. The trick is: if the header file defining `vtss_port_mask_t` is included by a C++ compiler, then it defines the `operator[]` without changing the API.

3.2.3.2. Other arrays

Creating bit masks to replace arrays, will only work for port-lists. Doing worst case allocation for other types will cause a significant overhead in terms of memory, which can not be accepted.

Instead of doing worst-case allocation, the data structures needs to be re-organized and/or the application need to dynamic allocate arrays based on the capability values returned by `mesa_capability`.

Lets consider an other example where the number of ports is being used to specify the length of an array. The following two snippets shows how a function in the Unified-API, and the MESA API.

`vtss_l2_api.h`

```
1 vtss_rc vtss_vlan_tx_tag_set(
2     const vtss_inst_t      inst,
3     const vtss_vid_t       vid,
4     const vtss_vlan_tx_tag_t tx_tag[VTSS_PORT_ARRAY_SIZE]);
```

`mscc/ethernet/switch/api/l2.h`

```
1 mesa_rc mesa_vlan_tx_tag_set(const mesa_inst_t      inst,
2                             const mesa_vid_t       vid,
3                             const uint32_t         count,
4                             const mesa_vlan_tx_tag_t *const tx_tag);
```

The Unified-API know the dimension at compile time, and can therefore use it directly in the prototype. This is not the case in the MEAS library, where it is parsed as a pointer, along with a `count`.

To call the MESA function from a C application will require that either the application know the size of the array, or that it uses the capability system to determinate it. A C++ application can benefit from using the `CapArray` to take care of the memory handling.

Lets have a look at the three different scenarios (all examples assume the API has initialized before use):

`c_application_with_static_sizes.c`

```
1 #define PORT_CNT 8
2 int some_function() {
3     mesa_vlan_tx_tag_t tags[PORT_CNT];
4     memset(tags, PORT_CNT, sizeof(mesa_vlan_tx_tag_t));
5     // fill in the data
6     return mesa_vlan_tx_tag_set(NULL, 1, PORT_CNT, tags);
7 }
```

c_application_with_dynamic_sizes.c

```
1 int some_function() {
2     int res = 0;
3     int cnt = mesa_capability(NULL, MESA_CAP_PORT_CNT);
4     mesa_vlan_tx_tag_t *tags = calloc(cnt, sizeof(mesa_vlan_tx_tag_t));
5     if (!tags) {
6         return -1;
7     }
8
9     // fill in the data
10
11     res = mesa_vlan_tx_tag_set(NULL, 1, PORT_CNT, tags);
12
13     free(tags);
14
15     return res;
16 }
```

caparray_application.cxx

```
1 int some_function() {
2     CapArray<mesa_vlan_tx_tag_t, MESA_CAP_PORT_CNT> tags;
3     // fill in the data
4     return mesa_vlan_tx_tag_set(NULL, 1, tags.size(), tags);
5 }
```

The three examples does the same thing, it is just different ways of doing it. Application should choose the approach that fits best into the current design, and the one they are most comfortable with.

3.2.4. Case study - QOS

Following is a more complete example taken from the QOS module, more specific the `vtss_qos_port_conf_t` structure. The structure has lots of `ifdef`, arrays with target specific constants, and features that takes up significant amount of space that is kind of unrelated. See the original code below:

vtss_qos_api.h

```

1  typedef struct {
2      vtss_policer_t policer_port[VTSS_PORT_POLICERS];
3      vtss_policer_ext_t policer_ext_port[VTSS_PORT_POLICERS];
4      vtss_policer_t policer_queue[VTSS_QUEUE_ARRAY_SIZE];
5      vtss_shaper_t shaper_port;
6      vtss_shaper_t shaper_queue[VTSS_QUEUE_ARRAY_SIZE];
7      #if defined(VTSS_FEATURE_QOS_EGRESS_QUEUE_SHAPERS_EB)
8          BOOL excess_enable[VTSS_QUEUE_ARRAY_SIZE];
9      #endif
10     #if defined(VTSS_FEATURE_QOS_EGRESS_QUEUE_SHAPERS_CRB)
11         BOOL credit_enable[VTSS_QUEUE_ARRAY_SIZE];
12     #endif
13     #if defined(VTSS_FEATURE_QOS_EGRESS_QUEUE_CUT_THROUGH)
14         BOOL cut_through_enable[VTSS_QUEUE_ARRAY_SIZE];
15     #endif
16     #if defined(VTSS_FEATURE_QOS_FRAME_PREEMPTION)
17         vtss_qos_fp_port_conf_t fp;
18         BOOL frame_preemption_enable[VTSS_QUEUE_ARRAY_SIZE];
19     #endif
20     vtss_prio_t default_prio;
21     vtss_tagprio_t usr_prio;
22     vtss_dp_level_t default_dpl;
23     vtss_dei_t default_dei;
24     BOOL tag_class_enable;
25     vtss_prio_t qos_class_map[VTSS_PCP_ARRAY_SIZE][VTSS_DEI_ARRAY_SIZE];
26     vtss_dp_level_t dp_level_map[VTSS_PCP_ARRAY_SIZE][VTSS_DEI_ARRAY_SIZE];
27     BOOL dscp_class_enable;
28     vtss_dscp_mode_t dscp_mode;
29     vtss_dscp_emode_t dscp_emode;
30     BOOL dscp_translate;
31     vtss_tag_remark_mode_t tag_remark_mode;
32     vtss_tagprio_t tag_default_pcp;
33     vtss_dei_t tag_default_dei;
34     vtss_tagprio_t tag_pcp_map[VTSS_PCP_ARRAY_SIZE][2];
35     vtss_dei_t tag_dei_map[VTSS_PCP_ARRAY_SIZE][2];
36     BOOL dwrr_enable;
37     #if defined(VTSS_FEATURE_QOS_SCHEDULER_DWRR_CNT)
38         u8 dwrr_cnt;
39     #endif
40     vtss_pct_t queue_pct[VTSS_QUEUE_ARRAY_SIZE];
41     #if defined(VTSS_FEATURE_QCL_DMAC_DIP)
42         BOOL dmac_dip;
43     #endif
44     #if defined(VTSS_FEATURE_QCL_KEY_TYPE)
45         vtss_vcap_key_type_t key_type;
46     #endif
47     #if defined(VTSS_FEATURE_QOS_WRED_V3)
48         vtss_wred_group_t wred_group;
49     #endif
50     #if (defined VTSS_FEATURE_QOS_COSID_CLASSIFICATION)
51         vtss_cosid_t cosid;
52     #endif
53     #if (defined VTSS_FEATURE_QOS_INGRESS_MAP)
54         vtss_qos_ingress_map_id_t ingress_map;

```

```

55 #endif
56 #if (defined VTSS_FEATURE_QOS_EGRESS_MAP)
57     vtss_qos_egress_map_id_t egress_map;
58 #endif
59 #if defined(VTSS_FEATURE_QOS_EGRESS_QUEUE_SHAPERS_TAS)
60     vtss_qos_qbv_port_conf_t qbv;
61 #endif
62 } vtss_qos_port_conf_t;

```

mssc/ethernet/switch/api/qos.h

```

1  typedef struct {
2      mesa_shaper_t          shaper;
3      mesa_prio_t           default_prio;
4      mesa_dp_level_t       default_dpl;
5      mesa_qos_port_tag_conf_t tag;
6      mesa_qos_port_dscp_conf_t dscp;
7      mesa_bool_t          dwrr_enable;
8      uint8_t              dwrr_cnt CAP(QOS_SCHEDULER_CNT_DWRR);
9      mesa_bool_t          dmac_dip CAP(QOS_QCL_DMACHIP);
10     mesa_vcap_key_type_t   key_type CAP(QOS_QCL_KEY_TYPE);
11     mesa_wred_group_t      wred_group CAP(QOS_WRED_GROUP_CNT);
12     mesa_cosid_t          cosid CAP(QOS_COSID_CLASSIFICATION);
13     mesa_qos_ingress_map_id_t ingress_map CAP(QOS_INGRESS_MAP_CNT);
14     mesa_qos_egress_map_id_t egress_map CAP(QOS_EGRESS_MAP_CNT);
15     mesa_qos_qbv_port_conf_t qbv CAP(QOS_EGRESS_QUEUE_SHAPERS_TAS);
16     mesa_qos_fp_port_conf_t fp CAP(QOS_FRAME_PREEMPTION);
17     mesa_qos_port_queue_conf_t queue[MESA_QUEUE_ARRAY_SIZE];
18 } mesa_qos_port_conf_t;
19
20 mesa_rc mesa_qos_port_conf_set(const mesa_inst_t inst,
21                               const mesa_port_no_t port_no,
22                               mesa_qos_port_conf_t *const conf);
23
24 typedef struct {
25     mesa_policer_t policer;
26     mesa_bool_t frame_rate;
27     mesa_dp_level_t dp_bypass_level CAP(QOS_PORT_POLICER_EXT_DPBL);
28     mesa_bool_t known_unicast CAP(QOS_PORT_POLICER_EXT_TTM);
29     mesa_bool_t known_multicast CAP(QOS_PORT_POLICER_EXT_TTM);
30     mesa_bool_t known_broadcast CAP(QOS_PORT_POLICER_EXT_TTM);
31     mesa_bool_t unknown_unicast CAP(QOS_PORT_POLICER_EXT_TTM);
32     mesa_bool_t unknown_multicast CAP(QOS_PORT_POLICER_EXT_TTM);
33     mesa_bool_t unknown_broadcast CAP(QOS_PORT_POLICER_EXT_TTM);
34     mesa_bool_t learning CAP(QOS_PORT_POLICER_EXT_TTM);
35     mesa_bool_t limit_noncpu_traffic CAP(QOS_PORT_POLICER_EXT_TTM);
36     mesa_bool_t limit_cpu_traffic CAP(QOS_PORT_POLICER_EXT_TTM);
37     mesa_bool_t flow_control;
38 } mesa_qos_port_policer_conf_t;
39
40 mesa_rc mesa_qos_port_police_conf_set(const mesa_inst_t inst,
41                                       mesa_port_no_t port_no, uint32_t cnt,
42                                       mesa_qos_port_policer_conf_t *const pol);

```

Following is a listing of the different techniques which has been applied in the example. Note that this example only covers the changes to `vtss_qos_port_conf_t`, similar initiatives will be need in almost all other structures defined in the `vtss_qos_api.h` header file.

1. Defines that is not active anymore has been removed. The QOS module has a long history, and earlier versions supported Luton28 and JR1 targets which is significant different from the current collection of chips. These old chips is not supported anymore, which means that many of the `ifdef` does not have any effect.
2. Members/variables that is only available on specific targets has been guarded with the `CAP()` macro.
3. Queue and police attributes were represented as chip dependent arrays in the current version, they have been separated out and is now handled by dedicated `get / set` functions that will allow run-time definitions of length parameter.

The result is an interface that is `ABI` compatible across the different targets, it does not cause any significant overhead and the changes in how the API must be accessed are minimal.

4. Build system, auto-generation and migration

The Unified-API and MESA library uses CMake as build system. The CMake system allows to specify what targets to include in a given build, and to use an existing cross compiler.

4.1. Development environment

Working with the source code raises some requirements to the development environment. This section will provide instructions on how to set-up a development machine based on `x86_64 Ubuntu 16.04LTS` installation. Other (recent) Linux distributions can be used, but that is not supported by MSCC. Setting up the development environment requires `root` access through the `sudo` command.

First step is to install a bunch of required packages using the package system provided by Ubuntu:

```
1 $ sudo apt-get install bc build-essential bzip2 coreutils cpio findutils gawk git
grep gzip libc6-i386 libcrypt-openssl-rsa-perl libncurses5-dev patch perl python
ruby sed squashfs-tools tcl tar wget libyaml-tiny-perl libcgi-fast-perl ruby-parslet
cmake
```

Next step is to download and install the binary BSP. This example will be using `2017.02-035` as example, but future releases may depend on newer versions.

Use the following link to browse the released BSPs: <http://mscc-ent-open-source.s3-website-eu-west-1.amazonaws.com/>. The steps below will download, install and test that the installed binaries work:

```
1 $ cd
2 $ wget -q http://mscc-ent-open-source.s3-eu-west-1.amazonaws.com/public_root/bsp/
mscc-brsdk-mips-2017.02-035.tar.gz
3 $ sudo mkdir -p /opt/mscc
4 $ sudo tar xf mscc-brsdk-mips-2017.02-035.tar.gz -C /opt/mscc
5 $ /opt/mscc/mscc-brsdk-mips-2017.02-035/stage2/smb/x86_64-linux/usr/bin/
mipsel-buildroot-linux-gnu-gcc --version
6 mipsel-linux-gcc.br_real (Buildroot 2016.05-git) 5.3.0
7 Copyright (C) 2015 Free Software Foundation, Inc.
8 This is free software; see the source for copying conditions. There is NO
9 warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
```

The final step is to extract the API sources. The example below shows how to extract API.

```
1 $ cd
2 $ tar -xf API_x.yy.tar.gz # File name/version may differ, depending on the release
```

4.2. Building the sources

Building sources using CMake, is a two-step operation. First step is to configure the project. To do that the `cmake` or the `ccmake` command is used. The latter provide a curses GUI, where the different options can be selected, while the first require you to specify all the options on the command line. For more details, read the manuals provided along with CMake.

When CMake has configured a project, it will generate a set of build files, in this case Makefile's, which can be used to do the actual build.

This example is configuring a project to use the compiler from the BSP which was just installed, and is enabling the ocelot demo target.



The toolchain (and thereby the cross compiler) must be specified on the command line the first time CMake is invoked. The only way to change the toolchain is to delete the build folder, and start over.

```

1  $ cd
2  $ cd API_X.Y/vtss_api    # Enter the source folder
3  $ mkdir build; cd build # Create and enter build directory
4  $ cmake -DCMAKE_TOOLCHAIN_FILE=/opt/mscc/mscc-brsdk-mips-2017.02-035/stage2/smb/
x86_64-linux/usr/share/buildroot/toolchainfile.cmake -Dmesa_demo_ocelot_vsc7514=ON ..
5  -- The C compiler identification is GNU 6.3.0
6  -- The CXX compiler identification is GNU 6.3.0
7  -- Check for working C compiler: /opt/mscc/mscc-brsdk-mips-2017.02-035/stage2/
smb/x86_64-linux/usr/bin/mipsel-buildroot-linux-gnu-gcc
8  -- Check for working C compiler: /opt/mscc/mscc-brsdk-mips-2017.02-035/stage2/
smb/x86_64-linux/usr/bin/mipsel-buildroot-linux-gnu-gcc - works
9  -- Detecting C compiler ABI info
10 -- Detecting C compiler ABI info - done
11 -- Detecting C compile features
12 -- Detecting C compile features - done
13 -- Check for working CXX compiler: /opt/mscc/mscc-brsdk-mips-2017.02-035/stage2/
smb/x86_64-linux/usr/bin/mipsel-buildroot-linux-gnu-g++
14 -- Check for working CXX compiler: /opt/mscc/mscc-brsdk-mips-2017.02-035/stage2/
smb/x86_64-linux/usr/bin/mipsel-buildroot-linux-gnu-g++ - works
15 -- Detecting CXX compiler ABI info
16 -- Detecting CXX compiler ABI info - done
17 -- Detecting CXX compile features
18 -- Detecting CXX compile features - done
19 -- Project name = vtss_api
20 -- Type = Release
21 -- cxx_flags = -D_LARGEFILE_SOURCE -D_LARGEFILE64_SOURCE
-D_FILE_OFFSET_BITS=64 -Os -std=c++11
22 -- c_flags = -D_LARGEFILE_SOURCE -D_LARGEFILE64_SOURCE
-D_FILE_OFFSET_BITS=64 -Os -Wall -Wno-array-bounds -fasynchronous-unwind-tables
-std=c99 -D_POSIX_C_SOURCE=200809L -D_BSD_SOURCE -D_DEFAULT_SOURCE
23 -- Looking for include file endian.h
24 -- Looking for include file endian.h - found
25 -- Looking for include file asm/byteorder.h
26 -- Looking for include file asm/byteorder.h - found
27 -- MESA layer: ON
28 -- Configuring done
29 -- Generating done
30 -- Build files have been written to: /home/anielsen/API_X.Y/mesa/build
31 $ make -j 10           # Build the sources
32 ...
33 $ ls lib*.so mesa/demo/*.mfi
34 libvsc7514_aqr.so mesa/demo/mesa_demo_ocelot_vsc7514.mfi

```

The build is producing a MESA library matching the vsc7514 target, and a demo images which can be used on the Ocelot reference boards.

4.2.1. Auto generated content

During the build-process, most of the MESA source code is auto-generated. The auto generator is parsing the header files (both Unified-API and MESA), and check if it can implement the body of the MESA functions automatic. If that is possible (because the differences follows simple rules), then it will do so. If not then it will skip implementing that function. Functions that can not be auto generated, needs to be generated manually.

The linker will ensure that the union of the manual implemented and the auto generated functions covers all functions in the MESA library.

The auto generated is found here (inside the build folder):

```
1 $ ls mesa-ag/  
2 include log.txt mesa mesa.cxx mesa.hxx mesa_pp_rename.h rename
```

The `mesa.cxx` and `mesa.hxx` is the auto generated part of the mesa implementation, it is just compiled into the MESA library as everything else. The `rename` and `mesa_pp_rename.h` is a complete list of symbols renamed in the MESA library, they are intended to be used for migrating applications - this is covered in the next section.

5. Application migration

Existing application written for the Unified API can migrated to using the MESA library. To ease that a list of symbol renames is provided by the MESA building system. MSCC has done the migration of the WebStaX product family, by using the same rename script in combination with the CapArrays. Depending on the size of the application, such a migration can be a considerable effort.



Customers are not forced to migrate to MESA, if they do not see any benefits in doing so, they are encourage to avoid it. MSCC has converted the WebStaX family to use the MESA layer, because it allows significant build time performance when building for many targets, and because it provides better separation of the application and API. The existing Unified API is still supported, and offers the same feature set and is supported in the same way.

The auto generator (`ag.rb`) is parsing the header files of the MESA library and the Unified API. This means that the `ag.rb` tool has a complete overview of all the symbols defined in the two header files. Having that overview, makes it easy to create a list of all the renames needed to migrate an application from using the Unified API to the MESA library.

The auto generator is creating a list, in two formats. One that can be used by `sed` to perform the renames in place (called `rename`), and another which can be included in a compilation unit (but in the top of a file) to set the pre-processor to do the rename (called `mesa_pp_rename.h`).

The following command can be used to do the rename in place, for a large number of files. (Remember to have a backup!):

```
1 $ cd /root/of/repository/to/perform/rename/on  
2 $ find . \( -name "*.h" -or -name "*.c" -or \  
3     -name "*.cxx" -or -name "*.c" \) \  
4     -exec sed -i /path/to/auto/generated/rename {} \;
```

6. Reference boards and demo system

The MESA project includes a simple demo application which is using the MESA and MEBA libraries. The application includes a number of modules, each implemented in a single C file. These modules include the following functionality:

- Main control and initialization of MEBA/MESA
- Port and PHY control
- MAC Address table
- VLAN control
- IP management via selected switch port
- Code trace
- Test commands
- Debug functions
- Command Line Interface (CLI)

The application binary (`mesa_demo`) supports a number of startup options, for example trace level control. It runs as a daemon, allowing the user to execute shell commands.

6.1. CLI Application

The CLI Application is simple client program, which communicates with the demo application using a socket interface. The CLI binary `cli` does the following:

- Reads an input command
- Sends the command to the demo application
- Writes the response to the terminal
- Exits

6.2. Building the demo application

The section called 'Building the sources' describes how to build the MESA library. This build job will also build the demo application, and organize it all into a `mfi` file that can be loaded on the corresponding reference board, if that has been enabled.

To expose what reference systems is supported, use the `ccmake .` command, and look for the `mesa_demo_*` settings. In this case `mesa_demo_ocelot_vsc7514` will be `ON`. Other components required by this demo application will also be `ON`.

The following sections assume that the `mfi` file is already build (in section 'Building the sources').

6.3. Download

The MFI file can now be loaded into RAM on a reference board. This requires the following:

- The reference board must have a Flash image, which can load the MFI file.

- The reference board must be setup with IP access to a TFTP server.
- The MFI file must be copied to the TFTP server.

The example below assumes that the MFI file has been copied to `new.img` in the TFTP home directory on a TFTP server with IP address 1.1.1.2. The switch IP address is set to 1.1.1.1/24 in the default VLAN. The configuration is stored as startup configuration, so that the system is ready to load the MFI file after reboot.

Note that the system is stopped in the bootloader (using `^C`) before the `ramload` command is given.

```
Username: admin
Password:
# conf term
config # int vlan 1
config-if-vlan # ip addr 1.1.1.1 255.255.255.0
config-if-vlan # end
# copy running-config startup-config
# platform debug allow
< Output from command not shown >
# debug firmware ramload tftp://1.1.1.2/new.img
< Output from load and boot process not shown >
^C
ramload
< Output from boot process and Linux startup not shown >
login: root
```

6.4. Running

After logging in as `root`, the demo application can be started and controlled using the CLI Application.


```

# mesa_demo -h
mesa_demo options:

-h                : Show this help text
-p <port>        : Enable IP management port
-t <module>:<group>:<level> : Set trace level for <module> and <group>, use '*' for
wildcard
# mesa_demo
# cli
Available Commands:

Help
Exit
IP Status
MAC Add <mac_addr> <port_list> [<vid>]
MAC Agetime [<age_time>]
MAC Delete <mac_addr> [<vid>]
MAC Dump
MAC Flush
MAC Lookup <mac_addr> [<vid>]
Port Flow Control [<port_list>] [enable|disable]
Port MaxFrame [<port_list>] [<max_frame>]
Port Mode [<port_list>] [10hdx|10fdx|100hdx|100fdx|1000fdx|auto]
Port State [<port_list>] [enable|disable]
Port Statistics [<port_list>] [clear|packets|bytes|errors|discards]
Test [<test_no>]
VLAN Add <vid> <port_list>
VLAN Delete <vid>
VLAN Filter [<port_list>] [enable|disable]
VLAN Frame [<port_list>] [all|tagged|untagged]
VLAN PVID [<port_list>] [<vid>]
VLAN Type [<port_list>] [unaware|c-port|s-port]
VLAN UVID [<port_list>] [all|none|pvid]
Debug API [<layer>] [<group>] [<port_list>] [full] [clear]
Debug Chip ID
Debug I2C Read <port_list> <i2c_addr> <addr> [<count>]
Debug I2C Write <port_list> <i2c_addr> <addr> <value>
Debug MMD Read <port_list> <mmd_list> <mmd_addr>
Debug MMD Write <port_list> <mmd_list> <mmd_addr> <value>
Debug PHY Read <port_list> <addr_list> [<page>]
Debug PHY Write <port_list> <addr_list> <value> [<page>]
Debug Port Polling [enable|disable]
Debug Sym Query <word128>
Debug Sym Read <word128>
Debug Sym Write <word128> <value32>
Debug Trace [<module>] [<group>] [off|error|info|debug|noise]
# cli port mode
Port  State      Mode      Flow Control  Rx Pause  Tx Pause  MaxFrame  Link
----  -
1     Enabled     Auto     Enabled      Disabled  Disabled  1518     1Gfdx
2     Enabled     Auto     Enabled      Disabled  Disabled  1518     1Gfdx
3     Enabled     1Gfdx    Disabled     Disabled  Disabled  1518     Down
4     Enabled     1Gfdx    Disabled     Disabled  Disabled  1518     1Gfdx
5     Enabled     1Gfdx    Disabled     Disabled  Disabled  1518     Down
6     Enabled     1Gfdx    Disabled     Disabled  Disabled  1518     Down

```

7	Enabled	10Gfdx	Disabled	Disabled	Disabled	1518	Down
8	Enabled	10Gfdx	Disabled	Disabled	Disabled	1518	Down
9	Enabled	Auto	Enabled	Disabled	Disabled	1518	1Gfdx

6.5. Adding a Test

The demo application includes a test module, which registers a CLI command `cli test` into the system. The command takes an integer argument identifying the test case to be executed. Additional test cases can be added in the `vtss_api/mesa/demo/test.c` file. The code below shows how the function `my_test` is added to the `test_table` in addition to existing ACL and EVC test cases. The test simply prints the number of packets received on the first switch port.

```
static mesa_rc my_test(void)
{
    mesa_port_counters_t counters;
    MESA_RC(mesa_port_counters_get(NULL, 0, &counters));
    cli_printf("Rx Packets: %llu\n", counters.rmon.rx_etherStatsPkts);
    return MESA_RC_OK;
}

static test_entry_t test_table[] = {
    {
        "ACL test",
        test_acl
    },
    {
        "EVC test",
        test_etc
    },
    {
        "My test",
        my_test
    },
};
```

After building, downloading and starting the demo application, the new command can be executed as shown below.

```
# cli test
Number  Description
-----  -
0      ACL test
1      EVC test
2      My test
# cli test 2
Rx Packets: 10
```

Appendix A: C++ CapArray

The CapArray is included as an example to the user (application developer), to show way of integrating the capabilities in the application. It is not a part of the MESA-API, and therefore not part of the UnifiedAPI/MESA support agreement.

The WebStaX family is using the CapArray heavily, and it is an official part of the WebStaX product, which is also supporting issues that may be caused by CapArray.



This is a C++ library, it depends on RAII, which means that it will not work (give wrong result, leak or crash) if combined with `memset`, `memcpy`, `memcmp` and other traditional C constructions. Do not use it, if you do not understand it!

The purpose of the CapArray, is mimic traditional C arrays as much as possible, but allow the array dimensions to be runtime determinate. It is not an alternative to `std::vector` as the size of the array is fixed (in oppesite to `std::vector` which is providing dynamic size).

The Current example supports up to 4 dimensional arrays, but can easily be extended further. Here is a small example using it:

`vtss_appl/main/caparray.hxx`

```
1  #include "caparray.hxx"
2  #include "mscc/ethernet/switch/api/capability.h"
3
4  struct Foo {
5      int a, b;
6  };
7
8  int main() {
9      CapArray<Foo, MESA_CAP_PORT_CNT> foo_array;
10
11     // Make sure that the API is initialized correctly before using.
12
13     for (size_t i = 0; i < foo_array.size(); i++) {
14         foo_array[i].a = i;
15         foo_array[i].b = 1;
16     }
17
18     CapArray<Foo, MESA_CAP_PORT_CNT, MESA_CAP_L3_RLEG_CNT> foo_array_array;
19     for (size_t i = 0; i < foo_array.size(); i++) {
20         for (size_t j = 0; j < foo_array[i].size(); j++) {
21             foo_array_array[i][j].a = i;
22             foo_array_array[i][j].b = j;
23         }
24     }
25
26 }
```

Following is the CapArray sources.

vtss_appl/main/caparray.hxx

```

1  /*
2  Copyright (c) 2004-2017 Microsemi Corporation "Microsemi".
3
4  Permission is hereby granted, free of charge, to any person obtaining a copy
5  of this software and associated documentation files (the "Software"), to deal
6  in the Software without restriction, including without limitation the rights
7  to use, copy, modify, merge, publish, distribute, sublicense, and/or sell
8  copies of the Software, and to permit persons to whom the Software is
9  furnished to do so, subject to the following conditions:
10
11  The above copyright notice and this permission notice shall be included in all
12  copies or substantial portions of the Software.
13
14  THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
15  IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
16  FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
17  AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
18  LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,
19  OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE
20  SOFTWARE.
21
22  */
23
24  template <typename T>
25  class A1 {
26  public:
27      A1(const A1 &) = delete;
28      A1 &operator=(const A1 &rhs) = delete;
29      A1(T *d, size_t _1) : data_(d), d1(_1) {}
30      A1(A1 &&rhs) : data_(rhs.data_), d1(rhs.d1) {}
31
32      // Needed to make it non-POD, for static analysis to catch ARRSZ problem
33      ~A1(){}
34
35      T &operator[](size_t idx) {
36          CAP_ARRAY_CHECK_DIM(idx, d1);
37          return data_[idx];
38      }
39      const T &operator[](size_t idx) const {
40          CAP_ARRAY_CHECK_DIM(idx, d1);
41          return data_[idx];
42      }
43      size_t size() const { return d1; }
44      size_t mem_size() const { return d1 * sizeof(T); }
45      const T *data() const { return data_; }
46      T *data() { return data_; }
47
48      template <class Q = T>
49      typename std::enable_if<std::is_pod<Q>::value, void>::type clear() {
50          memset(data_, 0, mem_size());
51      }
52
53  protected:
54      T *const data_;

```

```

55     const size_t d1;
56 };
57
58 template <typename T>
59 class A2 {
60     public:
61         A2(const A2 &) = delete;
62         A2 &operator=(const A2 &rhs) = delete;
63         A2(A2 &&rhs) : data_(rhs.data_), d1(rhs.d1), d2(rhs.d2) {}
64         A2(T *d, size_t _1, size_t _2) : data_(d), d1(_1), d2(_2) {}
65
66         // Needed to make it non-POD, for static analysis to catch ARR SZ problem
67         ~A2(){}
68
69         A1<T> operator[](size_t idx) {
70             CAP_ARRAY_CHECK_DIM(idx, d1);
71             return A1<T>(data_ + (idx * d2), d2);
72         }
73         const A1<T> operator[](size_t idx) const {
74             CAP_ARRAY_CHECK_DIM(idx, d1);
75             return A1<T>(data_ + (idx * d2), d2);
76         }
77         size_t size() const { return d1; }
78         size_t mem_size() const { return d1 * d2 * sizeof(T); }
79         const T *data() const { return data_; }
80         T *data() { return data_; }
81
82         template <class Q = T>
83         typename std::enable_if<std::is_pod<Q>::value, void>::type clear() {
84             memset(data_, 0, mem_size());
85         }
86
87     protected:
88         T *data_ = nullptr;
89         size_t d1 = 0, d2 = 0;
90 };
91
92
93 template <typename T>
94 class A3 {
95     public:
96         A3(const A3 &) = delete;
97         A3 &operator=(const A3 &rhs) = delete;
98         A3(A3 &&rhs) : data_(rhs.data_), d1(rhs.d1), d2(rhs.d2), d3(rhs.d3) {}
99         A3(T *d, size_t x, size_t y, size_t z) : data_(d), d1(x), d2(y), d3(z) {}
100
101         // Needed to make it non-POD, for static analysis to catch ARR SZ problem
102         ~A3(){}
103
104         A2<T> operator[](size_t idx) {
105             CAP_ARRAY_CHECK_DIM(idx, d1);
106             return A2<T>(data_ + (idx * d2 * d3), d2, d3);
107         }
108         const A2<T> operator[](size_t idx) const {
109             CAP_ARRAY_CHECK_DIM(idx, d1);

```

```

110     return A2<T>(data_ + (idx * d2 * d3), d2, d3);
111 }
112 size_t size() const { return d1; }
113 size_t mem_size() const { return d1 * d2 * d3 * sizeof(T); }
114 const T *data() const { return data_; }
115 T *data() { return data_; }
116
117 template <class Q = T>
118 typename std::enable_if<std::is_pod<Q>::value, void>::type clear() {
119     memset(data_, 0, mem_size());
120 }
121
122 protected:
123     T *data_ = nullptr;
124     size_t d1 = 0, d2 = 0, d3 = 0;
125 };
126
127 template <typename T, int... CAPS>
128 class CapArray;
129
130 template <typename T, int C1>
131 class CapArray<T, C1> {
132 public:
133     CapArray() {}
134     CapArray &operator=(const CapArray &rhs) {
135         init();
136         for (size_t i = 0; i < d1; ++i) data_[i] = rhs.data()[i];
137         return *this;
138     }
139
140     CapArray(const CapArray &rhs) {
141         init();
142         for (size_t i = 0; i < d1; ++i) data_[i] = rhs.data()[i];
143     }
144
145     template <class Q = T>
146     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator==(
147         const CapArray &rhs) const {
148         init();
149         return memcmp(data(), rhs.data(), mem_size()) == 0;
150     }
151
152     template <class Q = T>
153     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator!=(
154         const CapArray &rhs) const {
155         return !this->operator==(rhs);
156     }
157
158     void init() const {
159         if (data_) return;
160         d1 = mesa_capability(nullptr, C1);
161         data_ = (T *)VTSS_CALLOC_MODID(VTSS_MODULE_ID, d1, sizeof(T), __FILE__,
162             __LINE__);
163         if (!data_) return;
164         for (size_t i = 0; i < d1; ++i) new (&(data_[i])) T();

```

```

165     }
166
167     ~CapArray() {
168         if (!data_) return;
169         for (size_t i = 0; i < d1; ++i) data_[i].~T();
170         VTSS_FREE(data_);
171     }
172
173     T &operator[](size_t idx) {
174         init();
175         CAP_ARRAY_CHECK_DIM(idx, d1);
176         return data_[idx];
177     }
178
179     const T &operator[](size_t idx) const {
180         init();
181         CAP_ARRAY_CHECK_DIM(idx, d1);
182         return data_[idx];
183     }
184
185     size_t size() const {
186         init();
187         return d1;
188     }
189
190     size_t mem_size() const {
191         init();
192         return d1 * sizeof(T);
193     }
194
195     const T *data() const {
196         init();
197         return data_;
198     }
199
200     T *data() {
201         init();
202         return data_;
203     }
204
205     template <class Q = T>
206     typename std::enable_if<std::is_pod<Q>::value, void>::type clear() {
207         init();
208         memset(data_, 0, mem_size());
209     }
210
211     private:
212         mutable size_t d1;
213         mutable T *data_ = nullptr;
214 };
215
216 template <typename T, int C1, int C2>
217 class CapArray<T, C1, C2> {
218     public:
219         CapArray() {}

```

```

220     CapArray &operator=(const CapArray &rhs) {
221         init();
222         for (size_t i = 0; i < d1 * d2; ++i) data_[i] = rhs.data()[i];
223         return *this;
224     }
225
226     CapArray(const CapArray &rhs) {
227         init();
228         for (size_t i = 0; i < d1 * d2; ++i) data_[i] = rhs.data()[i];
229     }
230
231     template <class Q = T>
232     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator==(
233         const CapArray &rhs) const {
234         init();
235         return memcmp(data(), rhs.data(), mem_size()) == 0;
236     }
237
238     template <class Q = T>
239     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator!=(
240         const CapArray &rhs) const {
241         return !this->operator==(rhs);
242     }
243
244     void init() const {
245         if (data_) return;
246         d1 = mesa_capability(nullptr, C1);
247         d2 = mesa_capability(nullptr, C2);
248         data_ = (T *)VTSS_CALLOC_MODID(VTSS_MODULE_ID, d1 * d2, sizeof(T),
249             __FILE__, __LINE__);
250         if (!data_) return;
251         for (size_t i = 0; i < d1 * d2; ++i) new (&(data_[i])) T();
252     }
253
254     ~CapArray() {
255         if (!data_) return;
256         for (size_t i = 0; i < d1 * d2; ++i) data_[i].~T();
257         VTSS_FREE(data_);
258     }
259
260     A1<T> operator[](size_t idx) {
261         init();
262         CAP_ARRAY_CHECK_DIM(idx, d1);
263         return A1<T>(data_ + (idx * d2), d2);
264     }
265
266     const A1<const T> operator[](size_t idx) const {
267         init();
268         CAP_ARRAY_CHECK_DIM(idx, d1);
269         return A1<const T>(data_ + (idx * d2), d2);
270     }
271
272     size_t mem_size() const {
273         init();
274         return d1 * d2 * sizeof(T);

```



```

275     }
276
277     size_t size() const {
278         init();
279         return d1;
280     }
281
282     const T *data() const {
283         init();
284         return data_;
285     }
286
287     T *data() {
288         init();
289         return data_;
290     }
291
292     template <class Q = T>
293     typename std::enable_if<std::is_pod<Q>::value, void>::type clear() {
294         init();
295         memset(data_, 0, mem_size());
296     }
297
298     private:
299     mutable size_t d1, d2;
300     mutable T *data_ = nullptr;
301 };
302
303 template <typename T, int C1, int C2, int C3>
304 class CapArray<T, C1, C2, C3> {
305     public:
306     CapArray() {}
307     CapArray &operator=(const CapArray &rhs) {
308         init();
309         size_t n = d1 * d2 * d3;
310         for (size_t i = 0; i < n; ++i) data_[i] = rhs.data()[i];
311         return *this;
312     }
313
314     CapArray(const CapArray &rhs) {
315         init();
316         size_t n = d1 * d2 * d3;
317         for (size_t i = 0; i < n; ++i) data_[i] = rhs.data()[i];
318     }
319
320     template <class Q = T>
321     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator==(
322         const CapArray &rhs) const {
323         init();
324         return memcmp(data(), rhs.data(), mem_size()) == 0;
325     }
326
327     template <class Q = T>
328     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator!=(
329         const CapArray &rhs) const {

```

```

330     return !this->operator==(rhs);
331 }
332
333 void init() {
334     if (data_) return;
335     d1 = mesa_capability(nullptr, C1);
336     d2 = mesa_capability(nullptr, C2);
337     d3 = mesa_capability(nullptr, C3);
338     size_t n = d1 * d2 * d3;
339
340     data_ = (T *)VTSS_CALLOC_MODID(VTSS_MODULE_ID, n, sizeof(T), __FILE__,
341                                 __LINE__);
342     if (!data_) return;
343     for (size_t i = 0; i < n; ++i) new (&(data_[i])) T();
344 }
345
346 ~CapArray() {
347     if (!data_) return;
348     size_t n = d1 * d2 * d3;
349     for (size_t i = 0; i < n; ++i) data_[i].~T();
350     VTSS_FREE(data_);
351 }
352
353 A2<T> operator[](size_t idx) {
354     init();
355     CAP_ARRAY_CHECK_DIM(idx, d1);
356     return A2<T>(data_ + (idx * d2 * d3), d2, d3);
357 }
358
359 const A2<T> operator[](size_t idx) const {
360     init();
361     CAP_ARRAY_CHECK_DIM(idx, d1);
362     return A2<const T>(data_ + (idx * d2 * d3), d2, d3);
363 }
364
365 size_t mem_size() const {
366     init();
367     return d1 * d2 * d3 * sizeof(T);
368 }
369
370 size_t size() const {
371     init();
372     return d1;
373 }
374
375 const T *data() const {
376     init();
377     return data_;
378 }
379
380 T *data() {
381     init();
382     return data_;
383 }
384

```

```

385     template <class Q = T>
386     typename std::enable_if<std::is_pod<Q>::value, void>::type clear() {
387         init();
388         memset(data_, 0, mem_size());
389     }
390
391
392     private:
393         mutable size_t d1, d2, d3;
394         mutable T *data_ = nullptr;
395 };
396
397 template <typename T, int C1, int C2, int C3, int C4>
398 class CapArray<T, C1, C2, C3, C4> {
399     public:
400         CapArray() {}
401         CapArray &operator=(const CapArray &rhs) {
402             init();
403             size_t n = d1 * d2 * d3 * d4;
404             for (size_t i = 0; i < n; ++i) data_[i] = rhs.data()[i];
405             return *this;
406         }
407
408         CapArray(const CapArray &rhs) {
409             init();
410             size_t n = d1 * d2 * d3 * d4;
411             for (size_t i = 0; i < n; ++i) data_[i] = rhs.data()[i];
412         }
413
414     template <class Q = T>
415     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator==(
416         const CapArray &rhs) const {
417         init();
418         return memcmp(data(), rhs.data(), mem_size()) == 0;
419     }
420
421     template <class Q = T>
422     typename std::enable_if<std::is_pod<Q>::value, bool>::type operator!=(
423         const CapArray &rhs) const {
424         return !this->operator==(rhs);
425     }
426
427     void init() {
428         if (data_) return;
429         d1 = mesa_capability(nullptr, C1);
430         d2 = mesa_capability(nullptr, C2);
431         d3 = mesa_capability(nullptr, C3);
432         d4 = mesa_capability(nullptr, C4);
433         size_t n = d1 * d2 * d3 * d4;
434
435         data_ = (T *)VTSS_CALLOC_MODID(VTSS_MODULE_ID, n, sizeof(T), __FILE__,
436                                     __LINE__);
437         if (!data_) return;
438         for (size_t i = 0; i < n; ++i) new (&(data_[i])) T();
439     }

```

```
440
441 ~CapArray() {
442     if (!data_) return;
443     size_t n = d1 * d2 * d3 * d4;
444     for (size_t i = 0; i < n; ++i) data_[i].~T();
445     VTSS_FREE(data_);
446 }
447
448 A3<T> operator[](size_t idx) {
449     init();
450     CAP_ARRAY_CHECK_DIM(idx, d1);
451     return A3<T>(data_ + (idx * d2 * d3 * d4), d2, d3, d4);
452 }
453
454 const A3<const T> operator[](size_t idx) const {
455     init();
456     CAP_ARRAY_CHECK_DIM(idx, d1);
457     return A3<const T>(data_ + (idx * d2 * d3 * d4), d2, d3, d4);
458 }
459
460 size_t mem_size() const {
461     init();
462     return d1 * d2 * d3 * d4 * sizeof(T);
463 }
464
465 size_t size() const {
466     init();
467     return d1;
468 }
469
470 const T *data() const {
471     init();
472     return data_;
473 }
474
475 T *data() {
476     init();
477     return data_;
478 }
479
480 template <class Q = T>
481 typename std::enable_if<std::is_pod<Q>::value, void>::type clear() {
482     init();
483     memset(data_, 0, mem_size());
484 }
485
486
487 private:
488     mutable size_t d1, d2, d3, d4;
489     mutable T *data_ = nullptr;
490 };
```

References

- [AN1007] MSCC Ethernet API Software
- [ug1068] SW Introduction to WebStaX under Linux
- [ug1069] MEBA Programmers Guide