



endace
a c c e l e r a t e d

dsm_loader User Guide

EDM04-07



Protection Against Harmful Interference

When present on equipment this manual pertains to, the statement "This device complies with part 15 of the FCC rules" specifies the equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the Federal Communications Commission [FCC] Rules.

These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment.

This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications.

Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at their own expense.

Extra Components and Materials

The product that this manual pertains to may include extra components and materials that are not essential to its basic operation, but are necessary to ensure compliance to the product standards required by the United States Federal Communications Commission, and the European EMC Directive. Modification or removal of these components and/or materials, is liable to cause non compliance to these standards, and in doing so invalidate the user's right to operate this equipment in a Class A industrial environment.

Disclaimer

Whilst every effort has been made to ensure accuracy, neither Endace Technology Limited nor any employee of the company, shall be liable on any ground whatsoever to any party in respect of decisions or actions they may make as a result of using this information.

Endace Technology Limited has taken great effort to verify the accuracy of this manual, but nothing herein should be construed as a warranty and Endace shall not be liable for technical or editorial errors or omissions contained herein.

In accordance with the Endace Technology Limited policy of continuing development, the information contained herein is subject to change without notice.

Website

<http://www.endace.com>

Copyright 2008 Endace Technology Ltd. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the Endace Technology Limited.

Endace, the Endace logo, Endace Accelerated, DAG, NinjaBox and NinjaProbe are trademarks or registered trademarks in New Zealand, or other countries, of Endace Technology Limited. Applied Watch and the Applied Watch logo are registered trademarks of Applied Watch Technologies LLC in the USA. All other product or service names are the property of their respective owners. Product and company names used are for identification purposes only and such use does not imply any agreement between Endace and any named company, or any sponsorship or endorsement by any named company.

Use of the Endace products described in this document is subject to the Endace Terms of Trade and the Endace End User License Agreement (EULA).

Contents

Data Stream Manager (DSM) Introduction	1
<hr/>	
DSM Overview	1
Packet Filters	2
Load Balancing (Steering) Algorithms	3
Lookup Table	4
Output Record Format	4
Counters	4
Loading Configurations	5
<hr/>	
Card Initialization	5
Load DSM Firmware	5
Configure the Card	6
dsm_loader Overview	9
dsm_loader Usage	9
File Format Specification	11
<hr/>	
XML Element Specification	11
<dsm-config>	11
<filter>	12
<interface>	13
<steering>	13
<partial>	14
<stream>	14
<raw>	16
<sonet>	17
<ethernet> & <ethernet-vlan>	17
<hdlc-type>	18
<ethertype>	18
<mac-source> & <mac-dest>	18
<ipv4>	19
<ip-source> & <ip-dest>	20
<tcp> & <udp>	20
<source-port> & <dest-port>	21
<tcp-flags>	21
<icmp>	21
<icmp-code>	22
<icmp-type>	23
<partial-component>	23
<stream-component>	24
<vlan-id>	24
Example Filter Files	25
Simple Steering	25
Interface Steering	26
Simple Filtering	27
Complete Example	28
Appendix A – Raw DSM Configuration Output	31
<hr/>	
Filters	31
Lookup Table	32
Version History	33
<hr/>	

Data Stream Manager (DSM) Introduction

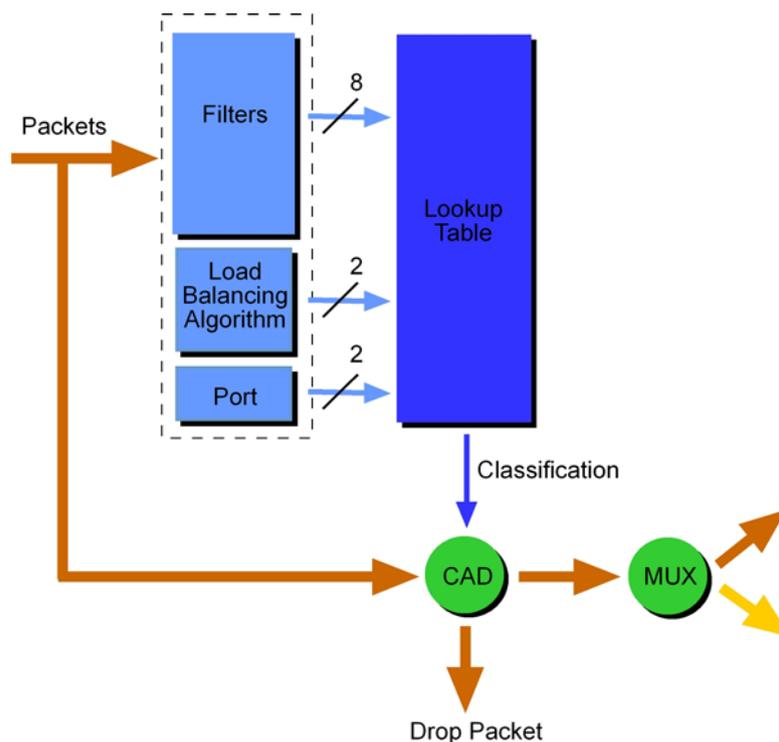
DSM Overview

The Data Stream Manager (DSM) is a feature supported on following DAG Cards:

- DAG 4.5G2
- DAG 4.5G4
- DAG 5.0SG2
- DAG 5.2X
- DAG 5.4S-12
- DAG 5.4S-48
- DAG 7.5G2
- DAG 7.5G4
- DAG 8.2X

DSM provides functionality to drop or route packets to a particular receive stream based on the packet contents, physical port and the output of two load balancing algorithms. The DSM logic is implemented in firmware on the DAG card, it does not require host CPU intervention once configured.

The following diagram shows the logical flow of packet records in the DSM module.



Packets are received from the line and stamped with an ERF (Endace Record Format) header, then passed along to the filter and load balancing block.

Filter / Load Balancing Block

The filter block applies eight bit-mask filters simultaneously to the first 64 bytes of the packet, producing a single true/false value for each filter. The load balancing (Load balancing (LB) is also known as Hash Load Balancing (HLB) or simply as just steering algorithm.) block applies two algorithms to the packet data, again producing one true/false boolean output per algorithm.

Lookup Table Block

Accepting the filter and load balancing outputs is the lookup table, it also receives the physical port the packet arrived on and calculates a classification for the packet. The classification is also known as the color of the packet.

Colorizer and Drop Block

The color is then passed onto the Colorizer And Drop (CAD) block that checks if the packet should be dropped, if not the color is inserted into the ERF record header of the packet and then the packet is passed along to the packet record multiplexer.

For detailed information on ERF record formats refer to the *EDM11-01 EFT Types* document which is available from the Support section of the Endace website at <http://www.endace.com>.

Packet Record Multiplexer (ERF MUX)

The ERF MUX looks at the color information contained in the packet record and determines which receive stream the packet record should be routed to.

Packet Filters

Prior to packets being presented to the DSM module, they are stamped with an ERF record header and possibly snapped to a particular length (set by the slen card configuration option). This is standard DAG card behavior, but should be taken into account when using the DSM firmware as it could affect filter output.

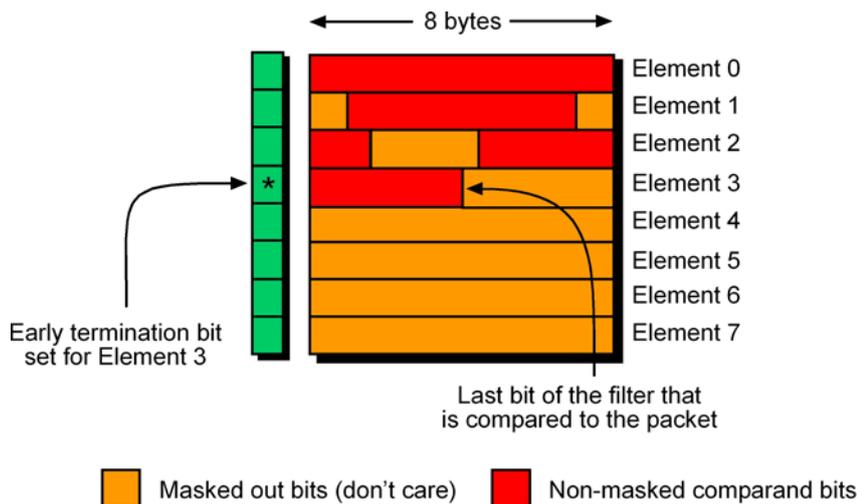
There are eight 64-byte bit masked filters inside the DSM module, each are compared against the packet in parallel. The first byte of the filter is compared against the first byte of the packet record after the ERF header, refer to the Endace Extensible Record Format document for more information on the packet record format. It is important to note that for Ethernet packets there are two bytes of padding added immediately after the ERF header, these padding bytes are the first to be compared against the filter.

Each filter outputs a Boolean `true` or `false` value that is provided to the lookup table for further classification.

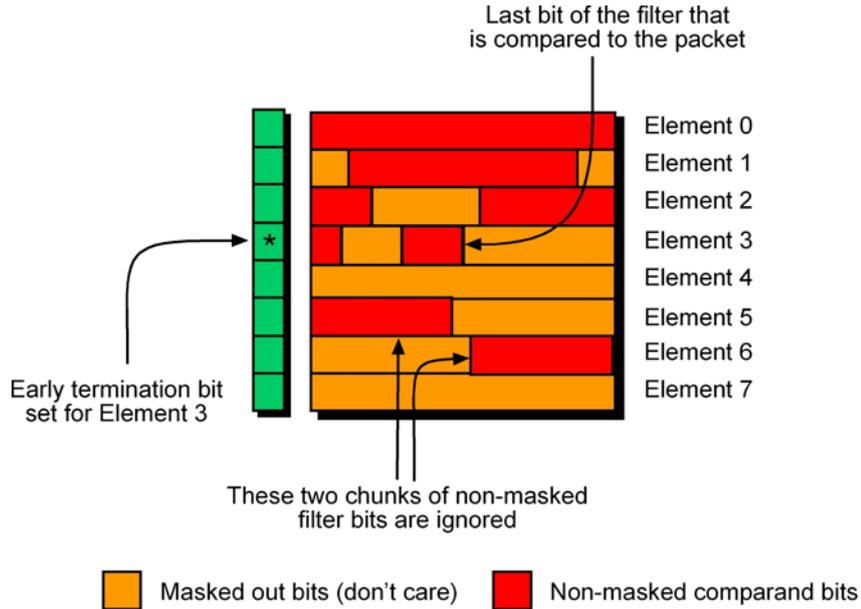
Filters also have an early termination option, this allows the user to specify on which 8-byte chunk (known as an element) of the filter contains the last byte to check. The early termination option is always specified on the last element in the filter (element 7). Packets that are smaller than the filter, as defined by the early termination option, always produce a false output regardless of the packet contents.

The following diagram shows a logical drawing of a filter, each of the rows represents 8 bytes of the filter (one element). In the diagram, the filter will be applied to the first 28 bytes (3 elements × 8 bytes + 4 non-masked bytes of element 3) of the packets rather than the full 64. Packets that are smaller than 28 bytes will produce a false output.

The following diagram shows an example of a filter with early termination.



The diagram below shows that any non-masked bytes of the filter that occur in elements after the early termination option are effectively ignored regardless of the packet length.



Load Balancing (Steering) Algorithms

Two load balancing algorithms are applied to the packet, each resulting in a Boolean output value, both outputs are provided to the lookup table for further classification. The first algorithm is a CRC calculation applied to the expected location of an IPv4 packet's source and destination address within the packet record. The second algorithm calculates the parity, across the expected location of an IPv4 packet's source and destination addresses.

For a random collection of packet data, both algorithms give an approximately 50:50 split of true and false outputs. The load balancing algorithms are fully implemented in firmware and are not user configurable.

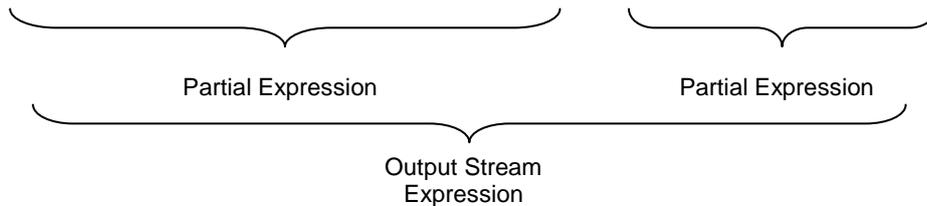
Lookup Table

The lookup table accepts the outputs from the filters, load balancing algorithm and the physical port number of the packet, to generate either a target stream number for the packet or a drop indication. The lookup table is fully user programmable, allowing for complex expressions to be constructed.

The DSM API provides a two stage implementation of the lookup table construction. The first stage involves creating one or more partial expressions, each parameter of the expression (or the inverse of the parameter) is logically OR'ed together to produce the partial expression. In the second stage, stream output expressions are constructed, containing one or more partial expressions, each partial expression (or the inverse of the partial expression) is AND'ed together.

Lookup Table Expression

Stream0 = (Filter0 OR Filter2 OR NOT Interface0) AND (Steering0 OR Filter6)



Packet records can be routed to only one stream, if more than one output expression returns a Boolean true value for a set of input parameters, the stream with the highest priority (lowest stream number) will receive the packet record. For example if the output stream expressions were the same for both stream 0 and stream 2, packet records that are accepted by the expression will only be routed to stream 0.

Output Record Format

Packets that are sent though the DSM are marked with a color value, this value encodes the outputs of the eight filters and two load balancing algorithms, as well as the target receive stream. Refer to Appendix A for the format of ERF record header including the color field.

Counters

The DSM module maintains thirteen counters, each counter is 32-bits and wraps back to zero on overflow.

DSM Counters

Type	Count	Description
Filter	8	Each filter has a counter indicating how many times the filter has output a true result.
Load Balancing	2	Both of the load balancing algorithms have a counter indicating how many true results have been generated.
Drop	1	Counts the number of packets that have been dropped.
Stream	n*	Each receive stream has a counter indicating the number of packet records that have been routed to that stream.

* the number of receive stream counters depends on the number of receive streams available on the card, currently this is 2.

Card Initialization

The DAG card may require configuration prior to loading the DSM settings. The following paragraphs illustrate the typical steps required for configuration.

Load DSM Firmware

By default DAG cards are not shipped with the required DSM firmware loaded into the FPGA, to load the correct firmware perform the following command.

```
dagrom -d0 -rvp -f <filename.bit>
```

Where `filename.bit` is the DSM FPGA image to load, and "0" is the number of the card. Depending on the type of card, it may be required to load the packet processing FPGA as well, refer to your DAG Card Guide for more information.

Configure the Card

Refer to the DAG Card User Guide for details on how to configure the card for your particular network settings. Usually the DAG card configuration is independent of the DSM configuration, however the snap length configuration attribute has a direct bearing on the DSM functionality.

Note: "1" indicates the condition is present on the link "0" indicates the condition is not present on the link.

DAG 4.5G2 & DAG 4.5G4 Card

Initialize the DAG 4.5G2 or DAG4.5G4 card using the following command:

```
dagconfig -d0 default
```

Verify the card has link and is initialized correctly using the following command:

```
dagconfig -d0 -si
```

After initializing the output should look as follows:

Port	Link	Plink	RFault	LOF	LOS
A	1	1	0	0	0
B	1	1	0	0	0

DAG 5.0SG2 Card

Initialize the DAG 5.0SG2 card using the following command:

```
dagconfig -d0 default
```

Verify the card has link and is initialized correctly using the following command:

```
dagconfig -d0 -si
```

After initializing the output should look as follows:

SONET

Port	lock	los	tx_lock_error	ais	b1	b2	b3	crc_error
A	1	0	0	0	0	0	0	0
B	1	0	0	0	0	0	0	0
fifo_empty	fifo_full	fifo_overflow	lof	lop	oof	rei_error	rdi_error	
0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	

Ethernet

Port	Lock	Los	Tx_lock_error
A	1	0	0
B	1	0	0

DAG 5.2X Card

Initialize the DAG 5.2X card using the following command:

```
dagconfig -d0 default
```

Verify the card has link and is initialized correctly using the following command:

```
dagconfig -d0 -si
```

After initializing the output should look as follows:

Port	Link	TXfault	RXFault	Lock
A	1	0	0	1
A	1	0	0	1

DAG 5.4S-12 Card

Initialize the DAG 5.4S-12 card using the following command:

```
dagconfig -d0 default
```

Verify the card has link and is initialized correctly using the following command:

```
dagconfig -d0 -si
```

After initializing the output should look as follows:

SONET

Port	lock	los	tx_lock_error	ais	b1	b2	b3	crc_error
A	1	0	0	0	0	0	0	0
B	1	0	0	0	0	0	0	0

fifo_empty	fifo_full	fifo_overflow	lof	lop	oof	rei_error	rdi_error
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

DAG 5.4S-48

Initialize the DAG 5.4S-48 card using the following command:

```
dagconfig -d0 default
```

Verify the card has link and is initialized correctly using the following command:

```
dagconfig -d0 -si
```

After initializing the output should look as follows:

SONET

Port	lock	los	tx_lock_error	ais	b1	b2	b3	crc_error
A	1	0	0	0	0	0	0	0
B	1	0	0	0	0	0	0	0

fifo_empty	fifo_full	fifo_overflow	lof	lop	oof	rei_error	rdi_error
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

Ethernet

Port	Lock	Los	Tx_lock_error
A	1	0	0
B	1	0	0

DAG 7.5G2 & DAG 7.5G4 Card

Initialize the DAG 7.5G2 or DAG 7.5G4 card using the following command:

```
dagconfig -d0 default
```

Verify the card has link and is initialized correctly using the following command:

```
dagconfig -d0 -si
```

After initializing the output should look as follows:

Port	pll_lock	lock	los	reset_done	link	lof
A	1	1	0	0	1	0
B	1	1	0	0	1	0

DAG 8.2X Card

Initialize the DAG 8.2X card using command

```
dagconfig -d0 default
```

Verify the card has link and is initialized correctly using command

```
dagconfig -d0 -si
```

After initializing the output should look as follows:

Port	link	fault	tx_fault	rx_fault	lock
A	1	0	0	0	1

Status Conditions

A definition of each of the status conditions is described below:

Condition	Description
Ais	Alarm indication signal. indicates a sonet/sdh remote aps error.
B1	A b1 parity error has occurred.
B2	A b2 parity error has occurred.
B3	A b3 parity error has occurred.
Crc_error	Indicates a crc error has occurred.
Fifo_empty	Indicates a fifo empty error has occurred.
Fifo_full	Indicates a fifo full error has occurred.
Fifo_overflow	Indicates a fifo overflow error has occurred.
Link	The link is fully validated
Lock	Card is locked to the incoming signal and can capture data.
Lof	Loss of frame. indicates oof has been asserted for more than 3 ms.
Lop	Loss of pointer
Los	Loss of Signal There is either no signal at the receiver or the optical signal strength is too low for the card to recognize.
Oof	Out of frame. indicates the section overhead processor is not locked to the sonet stream.
Plink	The peer link is up.
Port	Indicates the port the status conditions apply to.
Rdi_error	Remote Defect Indicator. When this condition exists the link will still carry data.
Rei_error	Remote error indicator.
RFault	There is an error at the remote end of the link
RxFault	There is a receive fault
Tx_lock_error	Lock error on transmit
TXFault	There is a transmission fault

dsm_loader Overview

The command line `dsm_loader` application is a tool that accepts a DSM configuration file and loads the details into the DSM firmware on supported DAG cards. It utilizes the DSM software API library to configure the DSM.

dsm_loader Usage

The following arguments can be used with `dsm_loader`

```
Usage : dsm_loader -d <device> [options] -f <config file>
```

Options	Description
-h, --help, -?, --usage	Display usage (help) information
-v, --verbose	increase verbosity
-d, --device <device>	DAG device to use
-f, --config_file <filename>	configuration file for the filters and expressions
-y, --dont_download	don't download the configuration to the card, if not specified the configuration is always downloaded to the card
-s, --stats <seconds>	display DSM statistics every <seconds>
-o, --output_file <filename>	output file to dump the constructed filters and lookup table

The use of a file containing configuration details is mandatory; the format of the file is given in the File Format Specification section of this document.

The `-y` don't download option is useful if you want to simply parse a configuration file and observe the raw DSM configuration using the `-o` option. The file format used for the output file is detailed in [Appendix A - Raw DSM Configuration Output](#) (page 31).

Statistics can be displayed periodically using the `-s` option, the statistics consist of the counters maintained by the DSM firmware.

File Format Specification

The filter loader file is written in xml format, it must satisfy the format by having a single root node `<dsm-config>`. Detailed in the following sections are the possible elements that may be used inside the root element.

XML Element Specification

<dsm-config>

This must be the root element of the xml document.

Attributes

Attribute	Use	Type	Default Value	Description
version	Mandatory	float	-	Defines the version of the file, currently the only version is 1.0. This document describes version 1.0 only.

Elements

Element	Occurrences		Description
	Min	Max	
filter	0	7	Filter element that contains the details of a particular filter.
interface	0	Unlimited	Interface element that defines a particular interface and labels it.
steering	0	Unlimited	Steering element that labels a hash load balancing algorithm used by the card.
partial	0	Unlimited	Partial element that contains a collection of components that make up a partial expression.
stream	0	Unlimited	Stream element that defines the output of a particular receive stream. There should only be one stream element per receive stream supported by the card.

<filter>

This element defines the construction of a filter, multiple filters can be defined but only one filter element per filter number is allowed.

Attributes

Attribute	Use	Type	Default Value	Description
enabled	Optional	boolean	true	Defines the filter as either enabled or disabled. By default the filter is enabled. Possible values for this attribute are true or false.

Elements

Element	Occurrences		Description
	Min	Max	
name	1	1	The name of the filter, this is a user defined name that is used when creating partial expressions.
number	1	1	The number of the filter, this should be value in the range of 0-6. It is an error to have two filters with the same number.
sonet	0	1	Defines the filter as a sonet filter, the filter will be created with the assumption that packets being received are encapsulated in PoS frames.
ethernet	0	1	Defines the filter as an ethernet filter, the filter is constructed with the assumption that packets will be encapsulated inside an ethernet frame.
ethernet-vlan	0	1	Defines the filter as an ethernet with VLAN tags, the filter is constructed with the assumption that packets will be encapsulated inside a ethernet frame with a four byte VLAN tag.
raw	0	1	Defines the filter as a raw filter, raw filters have the filter bits defined directly in the xml file.
early-term	0	1	Defines the early termination length, this is the last filter element (8 byte chunk) of the filter that will be used to compare against the packet. If this element is not specified it defaults to the maximum size of the filter (currently 8 elements). The packet record is compared up to the last non-masked byte of the element, packets that are smaller than the early termination length are dropped. Example: <code><early-term>1</early-term></code> Refer to section 2.1.1 for more information on the early termination option.

<interface>

Defines an interface that can be used inside a partial expression.

Elements

Element	Occurrences		Description
	Min	Max	
name	1	1	The name of the interface, this is a user defined name that is used when creating partial expressions.
number	1	1	The number of the interface, this should be value in the range of 0-3.

<steering>

Defines a steering algorithm that can be used inside a partial expression.

Elements

Element	Occurrences		Description
	Min	Max	
name	1	1	The name of the interface, this is a user defined name that is used when creating partial expressions.
algorithm	1	1	<p>Defines the algorithm used for steering. The two possible values are parity and crc32.</p> <p>Examples:</p> <pre><steering> <name>steer0</name> <algorithm>parity</algorithm> </steering> <steering> <name>steer1</name> <algorithm>crc32</algorithm> </steering></pre>

<partial>

Defines a partial expression, this element is expected to contain one or more <partial-component> elements, each defining either a filter, interface or steering algorithm to OR together to create a complete partial expression. For example the following XML snippet will produce:

Example 1:

```
partial0 = filter0 OR NOT filter1 OR interface0
<partial>
  <name>partial0</name>
  <partial-component>filter0</partial-component>
  <partial-component invert="true">filter1</partial-component>
  <partial-component>interface0</partial-component>
</partial>
```

Elements

Element	Occurrences		Description
	Min	Max	
name	1	1	The name of the partial expression, this is a user defined name that is used when creating a stream output expressions.
partial-component	0	Unlimited	Specifies the name of a partial component that makes up the partial expression.

<stream>

This element specifies the partial expressions that are AND'ed together to create a complete stream output expression.

If a stream element is not specified for a particular stream number, that stream is effectively disabled, no packets regardless of filters, steering algorithms or interface will ever be sent to that stream.

Elements

Element	Occurrences		Description
	Min	Max	
number	1	1	The receive stream number, this must be an even decimal number between 0 and the maximum number of receive streams.
stream-component	1	Unlimited	Specifies the name of a partial expression that makes up the stream output expression.

Note: packets can't be routed to multiple streams, therefore if the result of the filters, steering algorithm, interface and output expression applied by the DSM produces multiple stream outputs, the DSM library will adjust the output expression automatically to allow only the highest priority stream to receive the packet. The lower the stream number the higher the priority.

The following examples illustrate simple stream output expressions.

Example 1:

```
stream0 = partial0 AND partial1 AND NOT partial2
```

```
<stream>
  <number>0</number>
  <stream-component>partial0</stream-component>
  <stream-component>partial1</stream-component>
  <stream-component invert="true">partial2</stream-component>
</stream>
```

Example 2:

```
stream2 = (filter0 OR filter6) AND NOT (interface0 OR NOT steering0 OR filter1)
```

```
<partial>
  <name>partial0</name>
  <partial-component>filter0</partial-component>
  <partial-component>filter6</partial-component>
</partial>

<partial>
  <name>partial1</name>
  <partial-component>interface0</partial-component>
  <partial-component invert="true">steering0</partial-component>
  <partial-component>filter1</partial-component>
</partial>

<stream>
  <number>2</number>
  <stream-component>partial0</stream-component>
  <stream-component invert="true">partial1</stream-component>
</stream>
```

Example 3:

It is possible to construct a stream output expression that will never accept, as in the following example:

```
stream0 = (filter0) AND NOT (filter0)
```

```
<partial>
  <name>partial0</name>
  <partial-component>filter0</partial-component>
</partial>

<stream>
  <number>0</number>
  <stream-component>partial0</stream-component>
  <stream-component invert="true">partial0</stream-component>
</stream>
```

The `dsm_loader` program doesn't check for such situations, it is the user responsibility to supply correctly structured output expressions.

<raw>

This element defines the parent filter as being of type raw. This element can only have child elements of type <word>, each word represents 32-bits of the filter, they are parsed in top down order with the first word element corresponding to the first 32-bits of the filter.

Elements

Element	Occurrences		Description
	Min	Max	
word	0	16	Specifies the 32-bits to filter on, the characters in the contents of the element must be either '0', '1', or '-' for don't care.

It is not necessary to define all 16 words that make a complete filter, instead words that aren't specified are assumed to be don't cares, the following example illustrates three filters that are identical:

Example 1:

```

<filter>
  <name>filter0</name>
  <raw>
    <word>10101010101010101010101010101010----</word>
    <word>1010-----101010101010101010101010</word>
    <word>-----</word>
    <word>-----</word>
    <word>-----</word>
    <word>-----</word>
    <word>-----</word>
    <word>-----</word>
    <word>-----</word>
    <word>-----</word>
    <word>-----</word>
  </raw>
</filter>
<filter>
  <name>filter1</name>
  <raw>
    <word>10101010101010101010101010101010----</word>
    <word>1010-----101010101010101010101010</word>
    <word>-----</word>
    <word>-----</word>
  </raw>
</filter>
<filter>
  <name>filter2</name>
  <raw>
    <word>10101010101010101010101010101010----</word>
    <word>1010-----101010101010101010101010</word>
  </raw>
</filter>

```

<sonet>

This element defines the parent filter element as a sonet type filter, sonet filters expect PoS packets and therefore adjust the position of the layer3 and layer4 filter elements accordingly. Defining a filter as being of type sonet, doesn't force the filter to be that type, if the DSM library detects that the card being configured is not sonet (for example ethernet instead) the sonet specific child elements are ignored and the library automatically compensates for the correct layer 2 encapsulation.

Elements

Element	Occurrences		Description
	Min	Max	
hdlc-type	0	1	Defines the PoS (HDLC/PPP) header to filter on.
ipv4	0	1	Defines the IPv4 fields to filter on.

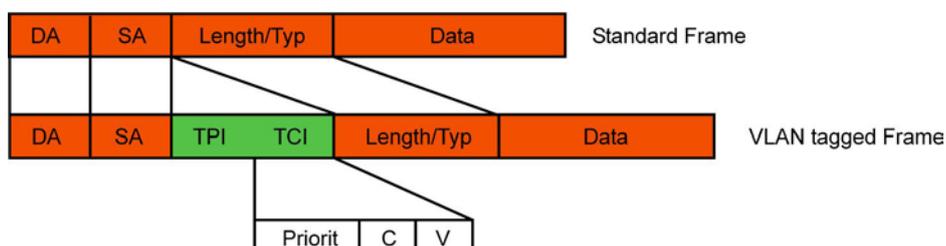
<ethernet> & <ethernet-vlan>

These elements defines the parent filter as a ethernet type filter (with or without VLAN), ethernet filters expect packets to be encapsulated in ethernet frames and therefore adjust the position of the layer3 and layer4 filter elements accordingly. As with the sonet element defining a filter as being an ethernet type doesn't force the filter to be that type, if the DSM library detects that the card being configured is not ethernet (for example sonet instead) the ethernet specific child elements are ignored and the library automatically compensates for the correct layer 2 encapsulation.

Elements

Element	Occurrences		Description
	Min	Max	
ethertype	0	1	Defines the 2 byte ethertype value to filter on.
mac-source	0	1	Defines the source ethernet MAC address to filter on.
mac-dest	0	1	Defines the destination ethernet MAC address to filter on.
ipv4	0	1	Defines the IPv4 fields to filter on.
vlan-id	0	1	Defines the 12-bit VLAN Id to filter on, this element is ignored if the <ethernet> element is the parent.

When <ethernet-vlan> is used, the DSM library will automatically add an entry in the filter for the VLAN ethertype (0x8100) and the user defined ethertype (if specified by the <ethertype> element) is relocated to the lower 16-bits of the VLAN tag as per the IEEE 802.1Q standard. The diagram below shows an example of an IEEE 802.1Q Standard, Tag-based VLAN.



<hdlc-type>

The contents of this element define the hdlc header to filter on, this is a 32-bit value. There is no mask for this element so the filter will hit on a direct match only.

Attributes

Attribute	Use	Type	Default Value	Description
hex	Optional	boolean	false	Indicates whether the contents of the element are in hexadecimal form or the default decimal form.

<ethertype>

The contents of this element define the ethertype value to filter on, this is a 16-bit value. There is no mask for this element so the filter will hit on a direct match only.

Attributes

Attribute	Use	Type	Default Value	Description
hex	Optional	boolean	false	Indicates whether the contents of the element are in hexadecimal form or the default decimal form.

<mac-source> & <mac-dest>

These elements defines the source/destination MAC address to filter on. If the <addr> child element is not present the address defaults to all zeros, if the <mask> child element is not present the mask defaults to all ones.

Elements

Element	Occurrences		Description
	Min	Max	
addr	0	1	<p>Defines the MAC address to filter on, the format of the contents should be X:X:X:X:X:X where X is either a decimal number (0-255) or a hexadecimal number (0-FF). The hex attribute determines the format of the address, by default decimal is used. Add hex="true" to the element to enable hex parsing.</p> <p>Examples: <addr>12:34:56:78:90:12</addr> <addr hex="true"> 99:AA:BB:CC:DD:EE:FF </addr></p>
mask	0	1	Defines the mask to use for the MAC address, the contents have the same format as the <addr> element.

<ipv4>

This element encapsulates the IPv4 specific filter elements. By defining this element it doesn't automatically change the ethertype/hdlc-type of the filter to match IPv4. If this is required you can do it manually by defining a <ethertype> or <hdlc-type> element inside the parent element.

Elements

Element	Occurrences		Description
	Min	Max	
ip-source	0	1	Defines the IPv4 source address to filter on.
ip-dest	0	1	Defines the IPv4 destination address to filter on.
no-ip-frags	0	1	If present, specifies that the filter will filter out all IPv4 packets that have been fragmented. If this element is not present the filter ignores whether or not the packet is a fragment.
ihl	0	1	Defines the IP header length to filter on, this element doesn't contain a mask therefore a direct match is required for a hit. The hex attribute can be specified if the contents of the element is in hexadecimal. The hex attribute defaults to "false". As with the IHL field in the IPv4 header this element should contain the length of the header in 32-bit words, for example <ihl>6</ihl> would define a header length of 24 bytes. The minimum header length is 5 and the maximum is 15. This element also effects the offset of any subsequent layer4 (UDP, TCP & ICMP) details,
tcp	0	1	Defines the layer4 type as TCP.
udp	0	1	Defines the layer4 type as UDP.
icmp	0	1	Defines the layer4 type as ICMP.

<ip-source> & <ip-dest>

These elements define the IPv4 source/destination addresses to filter on. If the <addr> child element is not present the address defaults to all zeros, if the <mask> child element is not present the mask defaults to all ones.

Elements

Element	Occurrences		Description
	Min	Max	
addr	0	1	<p>Defines the IPv4 address to filter on, the contents should be formatted like <code>x.x.x.x</code> where <code>x</code> is either a decimal number (0-255) or a hexadecimal number (0-FF). The <code>hex</code> attribute determines the format of the address, by default decimal is used. Add <code>hex="true"</code> to the element to enable hexadecimal parsing.</p> <p>Examples:</p> <pre><ip-source> <addr>192.168.0.0</addr> </ip-source> <ip-dest> <addr hex="true">C0.A8.0.0</addr> <mask hex="true">FF.FF.0.0</mask> </ip-dest></pre>
mask	0	1	<p>Defines the mask to use for the IPv4 address, the contents have the same format as the <addr> element.</p>

<tcp> & <udp>

These elements defines the layer4 type as either UDP or TCP. The protocol field in the IPv4 header is set to either `0x06` (TCP) or `0x11` (UDP) and the filter is updated.

Elements

Element	Occurrences		Description
	Min	Max	
source-port	0	1	Defines the UDP/TCP source port to filter on.
dest-port	0	1	Defines the UDP/TCP destination port to filter on.
tcp-flags	0	1	Defines the TCP flags to filter on, this element is only valid as the child of a <tcp> element,

<source-port> & <dest-port>

These elements define the UDP/TCP source/destination ports to filter on. If the <port> child element is not present the port defaults to all zeros, if the <mask> child element is not present the mask defaults to all ones.

Elements

Element	Occurrences		Description
	Min	Max	
port	0	1	<p>Defines the UDP/TCP port to filter on, the format of the contents should be either a decimal number (0-65535) or a hexadecimal number (0-FFFF). The hex attribute determines the format of the port, by default decimal is used. Add hex="true" to the element to enable hexadecimal port values.</p> <p>Examples:</p> <pre><dest-port> <port>80</port> <mask>255</mask> </dest-port> <source-port> <port hex="true">00AF</port> <mask hex="false">255</mask> </source-port></pre>
mask	0	1	<p>Defines the mask to use for the UDP/TCP port, the contents have the same format as the <port> element.</p>

<tcp-flags>

This element defines the TCP flags to filter on. If the <flags> element is not present the flags defaults to all zeros, if the <mask> element is not present the mask defaults to all ones.

Elements

Element	Occurrences		Description
	Min	Max	
flags	0	1	<p>Defines the TCP flags to filter on, the format of the contents should be either a decimal number (0-63) or a hexadecimal number (0-3F). The hex attribute determines the format of the flags, by default decimal is used. Add hex="true" to the element to enable hexadecimal port values.</p> <p>Examples:</p> <pre><flags>31</flags> < flags hex="true">1F</flags></pre>
mask	0	1	<p>Defines the mask to use for the TCP flags, the contents have the same format as the <flags> element.</p>

<icmp>

This element defines the layer4 type as ICMP. The protocol field in the IPv4 header is set to 0x01 (the ICMP protocol number) and the filter is updated.

Elements

Element	Occurrences		Description
	Min	Max	
icmp-code	0	1	Defines the ICMP code to filter on.
icmp-type	0	1	Defines the ICMP type to filter on.

<icmp-code>

This element defines the ICMP code to filter on. If the <code> element is not present the code defaults to all zeros, if the <mask> element is not present the mask defaults to all ones.

Elements

Element	Occurrences		Description
	Min	Max	
code	0	1	<p>Defines the ICMP code to filter on, the format of the contents should be either a decimal number (0-255) or a hexadecimal number (0-FF). The hex attribute determines the format of the code, by default decimal is used. Add hex="true" to the element to enable hexadecimal code values.</p> <p>Examples:</p> <pre><icmp-code> <code>31</code> <mask>255</mask> </icmp-code> <icmp-code> <code hex="true">1F</code> <mask hex="true">FF</mask> </icmp-code></pre>
mask	0	1	Defines the mask to use for the ICMP code, the contents have the same format as the <code> element.

<icmp-type>

This element defines the ICMP type to filter on. If the <type> element is not present the type defaults to all zeros, if the <mask> element is not present the mask defaults to all ones.

Elements

Element	Occurrences		Description
	Min	Max	
type	0	1	<p>Defines the ICMP type to filter on, the format of the contents should be either a decimal number (0-255) or a hexadecimal number (0-FF). The hex attribute determines the format of the type, by default decimal is used. Add hex="true" to the element to enable hexadecimal type values.</p> <p>Examples:</p> <pre><icmp-type> <type>31</type> <mask>255</mask> </icmp-type> <icmp-type> <type hex="true">1F</type> <mask hex="true">FF</mask> </icmp-type></pre>
mask	0	1	<p>Defines the mask to use for the ICMP type, the contents have the same format as the <type> element.</p>

<partial-component>

Defines either a filter, interface or steering algorithm that is part of a partial expression, each component is OR'ed together to create a partial expression. The component must correspond to a named <filter>, <interface> or <steering> element in the file.

Attributes

Attribute	Use	Type	Default Value	Description
invert	Optional	boolean	false	Indicates that the inversion of the component should be used in the partial expression.

Note: Although it is possible to specify a component and it's inverse in a partial expression, the actual outcome is that the last component is the one that is set. For example in the following xml snippet:

```
<partial>
  <name>partial0</name>
  <partial-component>filter0</partial-component>
  <partial-component invert="true">filter0</partial-component>
</partial>
```

you would expect the output to be `partial0 = filter0 OR NOT filter0` (an accept all expression), instead simply `partial0 = NOT filter0` will be produced, because the second <partial-component> element was processed last.

<stream-component>

Defines a partial expression that is part of a stream output expression, each partial expression is AND'ed together to form an output expression. The component must correspond to a named <partial> child element in the file.

Attributes

Attribute	Use	Type	Default Value	Description
invert	Optional	boolean	false	Indicates that the inversion of the partial should be used in the stream output expression.

<vlan-id>

This element defines the VLAN Id to filter on this is the 12-bit id contained within the VLAN tag of a packet. If the <id> child element is not present the ID defaults to all zeros, if the <mask> child element is not present the mask defaults to all ones.

Elements

Element	Occurrences		Description
	Min	Max	
id	0	1	<p>Defines the 12-bit VLAN ID to filter on, the format of the contents should be either a decimal number (0-4095) or a hexadecimal number (0-FFF). The hex attribute determines the format of the ID, by default decimal is used. Add hex="true" as an attribute to the element to enable hexadecimal ID values.</p> <p>Examples:</p> <pre><vlan-id> <id>31</id> <mask>63</mask> </vlan-id> <vlan-id> <id hex="true">1F</id> <mask hex="true">3F</mask> </vlan-id></pre>
mask	0	1	Defines the mask to use for the VLAN ID, the contents have the same format as the <id> element.

Example Filter Files

This section provides example DSM filter files.

Simple Steering

The following example uses the crc32 steering algorithm to approximately split the received packets into two streams (streams 0 and 2). The physical port and filters are ignored.

stream0 = steer

stream2 = NOT steer0

Example file:

```
<?xml version="1.0"?>

<dsm-config version="1.0">
  <steering>
    <name>steer0</name>
    <algorithm>crc32</algorithm>
  </steering>

  <partial>
    <name>partial0</name>
    <partial-component>steer0</partial-component>
  </partial>

  <stream>
    <number>0</number>
    <stream-component>partial0</stream-component>
  </stream>
  <stream>
    <number>2</number>
    <stream-component invert="true">partial0</stream-component>
  </stream>
</dsm-config>
```

Interface Steering

This example assumes that the configuration will be loaded into a DAG card that has four physical ports (interfaces). The configuration routes all traffic that arrives on ports 0 & 2 to receive stream 0, traffic received on port 1 is routed to stream 2 and any traffic that arrives on port 3 is dropped.

The outputs of both steering algorithms and the eight filters are ignored.

stream0 = iface0 OR iface2

stream2 = iface1

Example file:

```
<?xml version="1.0"?>

<dsm-config version="1.0">
  <interface>
    <name>iface0</name>
    <number>0</number>
  </interface>
  <interface>
    <name>iface1</name>
    <number>1</number>
  </interface>
  <interface>
    <name>iface2</name>
    <number>2</number>
  </interface>

  <partial>
    <name>partial0</name>
    <partial-component>iface0</partial-component>
    <partial-component>iface2</partial-component>
  </partial>
  <partial>
    <name>partial1</name>
    <partial-component>iface1</partial-component>
  </partial>

  <stream>
    <number>0</number>
    <stream-component>partial0</stream-component>
  </stream>
  <stream>
    <number>2</number>
    <stream-component>partial1</stream-component>
  </stream>
</dsm-config>
```

Simple Filtering

The following example routes all IPv4 packets that have a class B source IP address of 192.168.x.x to stream 0, all other TCP packets with a destination port of 80 are routed to stream 2 and everything else is dropped. The DAG card is assumed to be ethernet without VLAN.

Example file:

```
<?xml version="1.0"?>
<dsm-config version="1.0">

  <filter>
    <name>filter0</name>
    <number>0</number>
    <ethernet>
      <ipv4>
        <ip-source>
          <addr>192.168.0.0</addr>
          <mask hex="true">FF.FF.0.0</mask>
        </ip-source>
        <ip-dest>
          <addr hex="true">0.0.0.0</addr>
          <mask hex="true">0.0.0.0</mask>
        </ip-dest>
      </ipv4>
    </ethernet>
  </filter>
  <filter early-term="false">
    <name>filter1</name>
    <number>1</number>
    <ethernet>
      <ipv4>
        <tcp>
          <dest-port>
            <port>80</port>
            <mask hex="true">FFFF</mask>
          </dest-port>
        </tcp>
      </ipv4>
    </ethernet>
  </filter>
  <partial>
    <name>partial0</name>
    <partial-component>filter0</partial-component>
  </partial>
  <partial>
    <name>partial1</name>
    <partial-component>filter1</partial-component>
  </partial>

  <stream>
    <number>0</number>
    <stream-component>partial0</stream-component>
  </stream>
  <stream>
    <number>2</number>
    <stream-component>partial1</stream-component>
  </stream>

</dsm-config>
```

Complete Example

This example demonstrates how to create a complete example that utilises the filter, interface and steering algorithms to route the packets between the two receive streams.

stream0 = (filter0 OR filter1) AND NOT (iface0 OR iface2) AND (steer0)

stream2 = (filter2 OR filter3) AND NOT (iface1 OR iface3) AND (steer1)

Example file:

```
<?xml version="1.0"?>

<dsm-config version="1.0">

  <!-- ethernet IPv4 filter for odd source addresses -->
  <filter>
    <name>filter0</name>
    <number>0</number>
    <ethernet>
      <ipv4>
        <ip-source>
          <addr>0.0.0.1</addr>
          <mask>0.0.0.1</mask>
        </ip-source>
      </ipv4>
    </ethernet>
  </filter>

  <!-- ethernet IPv4 filter for odd destination addresses -->
  <filter>
    <name>filter1</name>
    <number>1</number>
    <ethernet>
      <ipv4>
        <ip-dest>
          <addr>0.0.0.1</addr>
          <mask>0.0.0.1</mask>
        </ip-dest>
      </ipv4>
    </ethernet>
  </filter>

  <!-- ethernet IPv4 filter for ICMP ping requests -->
  <filter>
    <name>filter2</name>
    <number>2</number>
    <ethernet>
      <ipv4>
        <icmp>
          <icmp-type>
            <type>8</type>
            <mask hex="true">FF</mask>
          </icmp-type>
        </icmp>
      </ipv4>
    </ethernet>
  </filter>
```

```
<!-- ethernet IPv4 filter for ICMP ping replies -->
<filter>
  <name>filter3</name>
  <number>3</number>
  <ethernet>
    <ipv4>
      <icmp>
        <icmp-type>
          <type>0</type>
          <mask hex="true">FF</mask>
        </icmp-type>
      </icmp>
    </ipv4>
  </ethernet>
</filter>

<!-- Interfaces -->
<interface>
  <name>iface0</name>
  <number>0</number>
</interface>
<interface>
  <name>iface1</name>
  <number>1</number>
</interface>
<interface>
  <name>iface2</name>
  <number>2</number>
</interface>
<interface>
  <name>iface3</name>
  <number>3</number>
</interface>

<!-- Steering algorithms -->
<steering>
  <name>steer0</name>
  <algorithm>crc32</algorithm>
</steering>
<steering>
  <name>steer1</name>
  <algorithm>parity</algorithm>
</steering>
```

```

<!-- Partial expressions -->
  <partial>
    <name>partial0</name>
    <partial-component>filter0</partial-component>
    <partial-component>filter1</partial-component>
  </partial>
  <partial>
    <name>partial1</name>
    <partial-component>filter2</partial-component>
    <partial-component>filter3</partial-component>
  </partial>
  <partial>
    <name>partial2</name>
    <partial-component>iface0</partial-component>
    <partial-component>iface2</partial-component>
  </partial>
  <partial>
    <name>partial3</name>
    <partial-component>iface1</partial-component>
    <partial-component>iface3</partial-component>
  </partial>
  <partial>
    <name>partial4</name>
    <partial-component>steer0</partial-component>
  </partial>
  <partial>
    <name>partial5</name>
    <partial-component>steer1</partial-component>
  </partial>

<!-- Output (stream) expressions -->
  <stream>
    <number>0</number>
    <stream-component>partial0</stream-component>
    <stream-component invert="true">partial2</stream-component>
    <stream-component>partial4</stream-component>
  </stream>
  <stream>
    <number>2</number>
    <stream-component>partial1</stream-component>
    <stream-component invert="true">partial3</stream-component>
    <stream-component>partial5</stream-component>
  </stream>

</dsm-config>

```

Appendix A – Raw DSM Configuration Output

This appendix describes the format of the `dsm_loader` generated raw configuration output files. The output files are generated when the `-o` option is used with `dsm_loader`.

The files are divided into two sections, the first containing the seven filter configurations, the second containing the lookup table.

Filters

The first line of the filter entry contains the attributes of the filter, possible values are given in the following table.

Filter Attributes

Name	Description
soft-filter	Indicates the filter is derived from a software configuration, rather than directly from the DAG card. This option is always specified.
virtual filter actual filter	Defines the virtual filter number and the actual filter number. The virtual filter number is the value used in the virtual configuration, whereas the actual filter number is the number of the filter inside the DAG card.
enabled	Indicates the filter is enabled
disabled	Indicates the filter is disabled
Early-terminate on element	Indicates the first element that has the early termination option set.

Following the attributes are sixteen lines of raw filter data, the first column contains the number of the filter, the next is the address of the filter data relative to the start of the filter, following that are 32 characters displaying the raw filter bits. A '-' signifies that bit is a 'don't care' value, '0' and '1' indicate bits that should be matched by the filter. The last two columns contain the hexadecimal value and mask words of the filter.

If a * is present at the end of the line it indicates that the early termination option has been set on that element, the early termination option is set over an element which is 64-bits in size, therefore the * will always be shown on two lines. By default the early termination option is always set on the last element of the filter if it hasn't been user defined.

Example 1. DSM filter output

```

soft-filter, virtual filter 2, actual filter 2, enabled, early-terminate on element 2
Filter2 : 00 [-----1010101010111011] Value[0000AABB] Mask[0000FFFF]
Filter2 : 04 [110011001101110111011101111111] Value[CCDDEEFF] Mask[FFFFFFFF]
Filter2 : 08 [00010010001101000101011001111000] Value[12345678] Mask[FFFFFFFF]
Filter2 : 0C [10010000001001010000001000000000] Value[90128100] Mask[FFFFFFFF]
Filter2 : 10 [----0000000100100100010001000100] Value[00124444] Mask[0FFFFFFF] *
Filter2 : 14 [-----] Value[00000000] Mask[00000000] *
Filter2 : 18 [-----] Value[00000000] Mask[00000000]
Filter2 : 1C [-----00000001-----] Value[00010000] Mask[00FF0000]
Filter2 : 20 [1100000010101000-----] Value[C0A80000] Mask[FFFF0000]
Filter2 : 24 [-----1010100000010111-----] Value[00A81700] Mask[00FFFF00]
Filter2 : 28 [1000000001010000-----] Value[80500000] Mask[FFFF0000]
Filter2 : 2C [-----] Value[00000000] Mask[00000000]
Filter2 : 30 [-----] Value[00000000] Mask[00000000]
Filter2 : 34 [-----] Value[00000000] Mask[00000000]
Filter2 : 38 [-----] Value[00000000] Mask[00000000]
Filter2 : 3C [-----] Value[00000000] Mask[00000000]

```

Lookup Table

The first four lines of the lookup table output, contain the statistics for the table. The `Bits per Entry` field, indicates how many bits are used to encoded the output of each entry, this is calculated automatically by the DSM API using the following formulae $\text{Bits per Entry} = \log_2(n) + 1$, where `n` is the number of receive memory holes. The `entries` field shows how many entries there are in the lookup table. The `Entries per Row` indicates the number of entries that can fit within a single row (16-bits), and `Rows` are the number of rows required to populate the complete table. The `Rows` referred to here is not related to a line of characters in the output file, rather it is a term used internally by the DSM API.

Each line of the output matches to a single entry in the lookup table, the `h1b-par`, `h1b-crc`, `iface` & `f7-f0` columns refer to the `true/false` input parameters into the lookup table. The `bm` column indicates either the output stream number or `drop`. The `row entry` column is the raw data that is programmed into the DAG card.

Example 2. DSM Lookup table output

```

Bits per Entry: 2
    Entries: 4096
Entries Per Row: 8
    Rows: 512

: h1b-par h1b-crc  iface  f7 f6 f5 f4 f3 f2 f1 f0 : bm  | row entry
-----
0x000 :    0    0    00    0 0 0 0 0 0 0 0 : drop | 0xAAAA
0x001 :    0    0    00    0 0 0 0 0 0 0 1 : drop | 0xAAAA
0x002 :    0    0    00    0 0 0 0 0 0 1 0 : drop | 0xAAAA
0x003 :    0    0    00    0 0 0 0 0 0 1 1 : drop | 0xAAAA
0x004 :    0    0    00    0 0 0 0 0 1 0 0 : drop | 0xAAAA
0x005 :    0    0    00    0 0 0 0 0 1 0 1 : drop | 0xAAAA
...
...
...
0xFFC :    1    1    11    1 1 1 1 1 1 0 0 : drop | 0x0202
0xFFD :    1    1    11    1 1 1 1 1 1 0 1 : 0    | 0x0202
0xFFE :    1    1    11    1 1 1 1 1 1 1 0 : 0    | 0x0202
0xFFF :    1    1    11    1 1 1 1 1 1 1 1 : 0    | 0x0202

```

Version History

Version	Date	Reason
1	March 2006	Initial revision
2	September 2007	New template and removal of some sections.
3	January 2008	Added information about 5.4S-12, 5.4S-48 and 8.1SX
4	February 2008	Removed 8.1SX reference.
5	November 2008	Updated front matter. Updated list of supported cards.

