



# SNF<sup>TM</sup> Application Programming Interface

Version 3.0.13.50838

January 5, 2018

All information contained in this document is proprietary to CSP, Inc. and may not be reproduced, distributed, or disseminated, in whole or in part, without the written permission of an authorized representative of CSP, Inc.

All specifications presented in this document are subject to change at any time, and without prior notice.

Myricom<sup>®</sup> and Myrinet<sup>®</sup> are registered trademarks of CSP, Inc. SNF<sup>TM</sup> is a trademark of CSP, Inc. Other trademarks appearing in this document are those of their respective owners.

©2008-2014, CSP, Inc.



# Contents

<b>1</b>	<b>SNF API Reference</b>	<b>1</b>
1.1	Sniffer Documentation	1
1.2	SNF API Reference	1
1.2.1	SNF API with Receive-Side Scaling	1
1.2.2	SNF API with Duplication	2
1.2.3	SNF API with Port Aggregation (Merging)	2
1.3	Examples	2
<b>2</b>	<b>Module Index</b>	<b>3</b>
2.1	API Reference	3
<b>3</b>	<b>Namespace Index</b>	<b>5</b>
3.1	Namespace List	5
<b>4</b>	<b>Data Structure Index</b>	<b>7</b>
4.1	Data Structures	7
<b>5</b>	<b>Module Documentation</b>	<b>9</b>
5.1	SNF API Reference	9
5.1.1	Detailed Description	11
5.1.2	API Reference	11
5.1.3	Processing Model	11
5.1.3.1	Multiple Receive Rings	12
5.1.3.2	Zero-Copy receives	12
5.1.3.3	API Software Overhead	12
5.1.3.4	Thread-Safety	12
5.1.4	Implementation details	13
5.1.4.1	Timestamps	13
5.1.4.2	Library Memory Consumption	13
5.1.4.3	Ethernet and Sniffer Driver Modes	13
5.1.5	Define Documentation	13
5.1.5.1	SNF_VERSION_API	13
5.1.6	Typedef Documentation	14
5.1.6.1	snf_handle_t	14
5.1.6.2	snf_ring_t	14
5.1.7	Enumeration Type Documentation	14
5.1.7.1	snf_link_state	14
5.1.7.2	snf_timesource_state	14
5.1.8	Function Documentation	14
5.1.8.1	snf_close	14

5.1.8.2	snf_freeifaddrs	15
5.1.8.3	snf_get_link_speed	15
5.1.8.4	snf_get_link_state	15
5.1.8.5	snf_get_timesource_state	16
5.1.8.6	snf_getifaddrs	16
5.1.8.7	snf_getportmask_linkup	16
5.1.8.8	snf_getportmask_valid	17
5.1.8.9	snf_init	17
5.1.8.10	snf_open	17
5.1.8.11	snf_open_defaults	18
5.1.8.12	snf_ring_close	19
5.1.8.13	snf_ring_getstats	19
5.1.8.14	snf_ring_open	19
5.1.8.15	snf_ring_open_id	20
5.1.8.16	snf_ring_portinfo	20
5.1.8.17	snf_ring_recv	21
5.1.8.18	snf_ring_recv_many	21
5.1.8.19	snf_ring_return_many	22
5.1.8.20	snf_set_app_id	23
5.1.8.21	snf_start	23
5.1.8.22	snf_stop	24
5.2	Receive-Side Scaling (RSS)	25
5.2.1	Detailed Description	25
5.2.2	Define Documentation	25
5.2.2.1	SNF_RSS_IPV4	25
5.2.3	Enumeration Type Documentation	25
5.2.3.1	snf_rss_mode_flags	25
5.2.3.2	snf_rss_params_mode	26
5.3	Open flags for process-sharing, port aggregation and packet duplication	27
5.3.1	Detailed Description	27
5.3.2	Define Documentation	27
5.3.2.1	SNF_F_AGGREGATE_PORTMASK	27
5.3.2.2	SNF_F_PSHARED	27
5.3.2.3	SNF_F_RX_DUPLICATE	27
5.4	Packet injection	28
5.4.1	Detailed Description	28
5.4.2	Typedef Documentation	28
5.4.2.1	snf_inject_t	28
5.4.3	Function Documentation	29
5.4.3.1	snf_get_injection_speed	29
5.4.3.2	snf_inject_close	29
5.4.3.3	snf_inject_getstats	29
5.4.3.4	snf_inject_open	30
5.4.3.5	snf_inject_sched	30
5.4.3.6	snf_inject_sched_v	31
5.4.3.7	snf_inject_send	32
5.4.3.8	snf_inject_send_v	33
5.5	Packet reflect to netdev (kernel stack)	35
5.5.1	Detailed Description	35
5.5.2	Typedef Documentation	35

5.5.2.1	snf_netdev_reflect_t . . . . .	35
5.5.3	Function Documentation . . . . .	35
5.5.3.1	snf_netdev_reflect . . . . .	35
5.5.3.2	snf_netdev_reflect_enable . . . . .	36
<b>6</b>	<b>Namespace Documentation</b>	<b>37</b>
6.1	snf Namespace Reference . . . . .	37
6.1.1	Detailed Description . . . . .	37
<b>7</b>	<b>Data Structure Documentation</b>	<b>39</b>
7.1	snf_ifaddrs Struct Reference . . . . .	39
7.1.1	Detailed Description . . . . .	39
7.1.2	Field Documentation . . . . .	39
7.1.2.1	pad . . . . .	39
7.1.2.2	snf_ifa_link_speed . . . . .	39
7.1.2.3	snf_ifa_link_state . . . . .	40
7.1.2.4	snf_ifa_macaddr . . . . .	40
7.1.2.5	snf_ifa_maxinject . . . . .	40
7.1.2.6	snf_ifa_maxrings . . . . .	40
7.1.2.7	snf_ifa_name . . . . .	40
7.1.2.8	snf_ifa_next . . . . .	40
7.1.2.9	snf_ifa_portnum . . . . .	40
7.2	snf_inject_stats Struct Reference . . . . .	40
7.2.1	Detailed Description . . . . .	40
7.2.2	Field Documentation . . . . .	41
7.2.2.1	inj_pkt_send . . . . .	41
7.2.2.2	nic_bytes_send . . . . .	41
7.2.2.3	nic_pkt_send . . . . .	41
7.3	snf_pkt_fragment Struct Reference . . . . .	41
7.3.1	Detailed Description . . . . .	41
7.3.2	Field Documentation . . . . .	41
7.3.2.1	length . . . . .	41
7.3.2.2	ptr . . . . .	41
7.4	snf_recv_req Struct Reference . . . . .	41
7.4.1	Detailed Description . . . . .	42
7.4.2	Field Documentation . . . . .	42
7.4.2.1	hw_hash . . . . .	42
7.4.2.2	length . . . . .	42
7.4.2.3	length_data . . . . .	42
7.4.2.4	pkt_addr . . . . .	42
7.4.2.5	portnum . . . . .	42
7.4.2.6	timestamp . . . . .	42
7.5	snf_ring_portinfo Struct Reference . . . . .	42
7.5.1	Detailed Description . . . . .	43
7.5.2	Field Documentation . . . . .	43
7.5.2.1	data_addr . . . . .	43
7.5.2.2	data_size . . . . .	43
7.5.2.3	portcnt . . . . .	43
7.5.2.4	portmask . . . . .	43
7.5.2.5	q_size . . . . .	43
7.5.2.6	ring . . . . .	43

7.6	snf_ring_qinfo Struct Reference	43
7.6.1	Detailed Description	44
7.6.2	Field Documentation	44
7.6.2.1	q_avail	44
7.6.2.2	q_borrowed	44
7.6.2.3	q_free	44
7.7	snf_ring_stats Struct Reference	44
7.7.1	Detailed Description	44
7.7.2	Field Documentation	44
7.7.2.1	nic_bytes_recv	44
7.7.2.2	nic_pkt_bad	45
7.7.2.3	nic_pkt_dropped	45
7.7.2.4	nic_pkt_overflow	45
7.7.2.5	nic_pkt_recv	45
7.7.2.6	ring_pkt_overflow	45
7.7.2.7	ring_pkt_recv	45
7.7.2.8	snf_pkt_overflow	45
7.8	snf_rss_mode_function Struct Reference	45
7.8.1	Detailed Description	45
7.8.2	Field Documentation	46
7.8.2.1	rss_context	46
7.8.2.2	rss_hash_fn	46
7.9	snf_rss_params Struct Reference	47
7.9.1	Detailed Description	47
7.9.2	Field Documentation	47
7.9.2.1	mode	47
7.9.2.2	params	47
7.9.2.3	rss_flags	47
7.9.2.4	rss_function	48

# Chapter 1

## SNF API Reference

### 1.1 Sniffer Documentation

The Sniffer User Guide covers NIC and software installation as well as instructions on how to use Sniffer with libpcap-based applications. It is available from <https://cspi.force.com/customersupport/>

### 1.2 SNF API Reference

The SNF API is available for applications that require tighter integration than libpcap with Sniffer. When used in its simplest form, the library resembles Libpcap in that the implementation expects a single thread to make successive calls to a receive function ([snf\\_ring\\_rcv](#)) to obtain the next available packet. Under a more advanced form, Sniffer implements a variation of the Receive-Side Scaling (RSS) feature that is present in some 10-Gigabit Ethernet drivers. However, Sniffer takes the additional step of implementing the RSS feature as multiple user-level zero-copy receive rings. Making the rings available in userspace provides two important advantages over all existing kernel-based packet capture solutions:

- No system calls are required to retrieve packets from a kernel-level queue, and users have zero-copy access to incoming packets;
- There is no need to funnel all incoming packets through a single Libpcap-like device as would be the case in kernel-based methods, even if RSS rings are allocated in the kernel.

#### 1.2.1 SNF API with Receive-Side Scaling

Sniffer translates the RSS feature into multiple rings in that data is hashed across many receive rings (or buffers or slices as is referenced in the myri10ge documentation). This feature assumes that users maintain a 1-to-1 relationship between user threads and rings. With each new call to [snf\\_ring\\_rcv](#), it is assumed that the previous packet in the ring has been completely consumed.

By default, the Sniffer implementation uses a deterministic hashing function to make sure that packets that are contained in a particular TCP or UDP flow are always delivered to the same ring (and hence to the same analysis thread). This hashing function resembles the hashing mechanisms used in existing RSS drivers.

### 1.2.2 SNF API with Duplication

While multiple rings are primarily designed to partition the incoming packet capture across multiple capture consuming rings, it is also possible to force each received packet to be duplicated into each ring such that every consuming ring obtains its own copy of every incoming packet. The duplication is handled by the Sniffer software on the host where there is typically plenty of memory bandwidth compared to the PCIe bus. Packet duplication can be enabled by setting the [SNF\\_F\\_RX\\_DUPLICATE](#) flag in [snf\\_open](#).

### 1.2.3 SNF API with Port Aggregation (Merging)

With Sniffer 2.0, it is now possible to logically aggregate packets from two or more Ethernet ports. The functionality can be extended through to consumers that employ RSS or duplication. This feature can be enabled in [snf\\_open](#) by setting the [SNF\\_F\\_AGGREGATE\\_PORTMASK](#) flag and passing a bitmask of ports to aggregate in the **portnum** parameter. As a convenience, functions are also available to return portmasks for valid ports ([snf\\_getportmask\\_valid](#)) and active ports ([snf\\_getportmask\\_linkup](#)). As a result of calling [snf\\_ring\\_recv](#), packets from one or more ports will be received.

## 1.3 Examples

Tests are available from **bin/tests** of the install directory in binary form and in **share/doc/examples** in source form. These tests mostly show different aspects of the SNF API and how to use its features.

**snf\_simple\_recv.c:** Simplest example of how to receive packets

**snf\_multi\_recv.c:** How to receive packets with multiple rings

**snf\_bridge.c:** Example of how to use SNF to create a transparent bridge to analyze traffic on one device and replay it on another

**snf\_pktgen.c:** How to generate packets for injection

**snf\_replay.c:** Example that uses SNF-level injection to replay a .pcap capture file

**snf\_basic\_diags:** Basic internal diagnostics which can be useful to verify that everything works as expected (source code not available)



## Chapter 2

# Module Index

### 2.1 API Reference

Here is a list of all modules:

SNF API Reference . . . . .	9
Receive-Side Scaling (RSS) . . . . .	25
Open flags for process-sharing, port aggregation and packet duplication . . . . .	27
Packet injection . . . . .	28
Packet reflect to netdev (kernel stack) . . . . .	35



## Chapter 3

# Namespace Index

### 3.1 Namespace List

Here is a list of all documented namespaces with brief descriptions:

<a href="#">snf</a> . . . . .	37
-------------------------------	----



## Chapter 4

# Data Structure Index

### 4.1 Data Structures

Here are the data structures with brief descriptions:

<a href="#">snf_ifaddrs</a>	39
<a href="#">snf_inject_stats</a>	40
<a href="#">snf_pkt_fragment</a>	
Fragment for snf_inject_send_v	41
<a href="#">snf_recv_req</a>	41
<a href="#">snf_ring_portinfo</a>	42
<a href="#">snf_ring_qinfo</a>	43
<a href="#">snf_ring_stats</a>	44
<a href="#">snf_rss_mode_function</a>	45
<a href="#">snf_rss_params</a>	47



## Chapter 5

# Module Documentation

### 5.1 SNF API Reference

API Reference for SNF and usage model.

#### Data Structures

- struct [snf\\_ifaddrs](#)
- struct [snf\\_recv\\_req](#)
- struct [snf\\_ring\\_qinfo](#)
- struct [snf\\_ring\\_portinfo](#)
- struct [snf\\_ring\\_stats](#)

#### Modules

- [Receive-Side Scaling \(RSS\)](#)
- [Open flags for process-sharing, port aggregation and packet duplication](#)
- [Packet injection](#)
- [Packet reflect to netdev \(kernel stack\)](#)

#### Defines

- `#define` [SNF\\_VERSION\\_API](#) 8

#### Typedefs

- typedef struct snf\_handle \* [snf\\_handle\\_t](#)  
*opaque snf device handle*
- typedef struct snf\_ring \* [snf\\_ring\\_t](#)  
*opaque snf ring handle*

## Enumerations

- enum `snf_link_state` { `SNF_LINK_DOWN` = 0, `SNF_LINK_UP` = 1 }
- enum `snf_timesource_state` {  
`SNF_TIMESOURCE_LOCAL` = 0, `SNF_TIMESOURCE_EXT_UNSYNCED`, `SNF_TIMESOURCE_EXT_SYNCED`, `SNF_TIMESOURCE_EXT_FAILED`,  
`SNF_TIMESOURCE_ARISTA_ACTIVE`, `SNF_TIMESOURCE_PPS` }

## Functions

- int `snf_init` (uint16\_t api\_version)  
Initialize Sniffer Library with api\_version == `SNF_VERSION_API`.
- int `snf_set_app_id` (int32\_t id)  
Set the application ID.
- int `snf_getifaddrs` (struct `snf_ifaddrs` \*\*ifaddrs\_o)
- void `snf_freeifaddrs` (struct `snf_ifaddrs` \*ifaddrs)
- int `snf_getportmask_valid` (uint32\_t \*mask\_o, int \*cnt\_o)
- int `snf_getportmask_linkup` (uint32\_t \*mask\_o, int \*cnt\_o)
- int `snf_open` (uint32\_t portnum, int num\_rings, const struct `snf_rss_params` \*rss\_params, int64\_t dataring\_sz, int flags, `snf_handle_t` \*devhandle)  
Open device for single or multi-ring operation.
- int `snf_open_defaults` (uint32\_t portnum, `snf_handle_t` \*devhandle)  
Open device for single or multi-ring operation.
- int `snf_start` (`snf_handle_t` devhandle)
- int `snf_stop` (`snf_handle_t` devhandle)
- int `snf_get_link_state` (`snf_handle_t` devhandle, enum `snf_link_state` \*state)
- int `snf_get_timesource_state` (`snf_handle_t` devhandle, enum `snf_timesource_state` \*state)
- int `snf_get_link_speed` (`snf_handle_t` devhandle, uint64\_t \*speed)
- int `snf_close` (`snf_handle_t` devhandle)  
Close port.
- int `snf_ring_open` (`snf_handle_t` devhandle, `snf_ring_t` \*ringh)
- int `snf_ring_open_id` (`snf_handle_t` devhandle, int ring\_id, `snf_ring_t` \*ringh)
- int `snf_ring_close` (`snf_ring_t` ringh)
- int `snf_ring_recv` (`snf_ring_t` ringh, int timeout\_ms, struct `snf_recv_req` \*recv\_req)  
Receive next packet from a receive ring.
- int `snf_ring_portinfo_count` (`snf_ring_t` ring, int \*count)  
For aggregated rings, return the number of physical subrings. If the ring is not aggregated, the count is set to 1.
- int `snf_ring_portinfo` (`snf_ring_t` ring, struct `snf_ring_portinfo` \*portinfo)  
Returns information for the ring. For aggregated rings, returns information for each of the physical rings. It is up to the user to make sure they have allocated enough memory to hold the information for all the physical rings in an aggregated ring.
- int `snf_ring_recv_qinfo` (`snf_ring_t` ring, struct `snf_ring_qinfo` \*)  
Return queue information from ring.
- int `snf_ring_recv_many` (`snf_ring_t` ring, int timeout\_ms, struct `snf_recv_req` \*req\_vector, int nreq\_in, int \*nreq\_out, struct `snf_ring_qinfo` \*qinfo)  
Receive and borrow many packets at once.



- int [snf\\_ring\\_return\\_many](#) (snf\_ring\_t ring, uint32\_t data\_qlen, struct [snf\\_ring\\_qinfo](#) \*qinfo)  
*Return packet space to receive ring.*
- int [snf\\_ring\\_getstats](#) (snf\_ring\_t ringh, struct [snf\\_ring\\_stats](#) \*stats)  
*Get statistics from a receive ring.*

### 5.1.1 Detailed Description

API Reference for SNF and usage model.

### 5.1.2 API Reference

The Sniffer API model is best summarized in the following steps:

- Initialize the API with the provided SNF\_VERSION\_API pre-define to ensure future ABI compatibility [snf\\_init](#).
- Query available Sniffer-capable devices if required ([snf\\_getifaddrs](#)) and open the device by port number ([snf\\_open](#)) with one or more rings.
- Have the current (or many other threads) open a receive ring for capture ([snf\\_ring\\_open](#)), then start packet capture ([snf\\_start](#)).
- Each ring is serviced in a loop by one or more consumers ([snf\\_ring\\_rcv](#)).
- Each ring is closed and finally, the device can be closed ([snf\\_ring\\_close](#) and [snf\\_close](#)).
- Packets can be injected by opening injection handles, which can be opened independently from the receive capability. Injection handles can be used to re-inject packets or as a packet generator.

### 5.1.3 Processing Model

The API promotes a processing model mindful of the multiple cores available on today's systems by enabling packets to be received through multiple receive rings. This feature appears in many 10Gbit Ethernet drivers as a Receive-Side Scaling (RSS) but instead of being an feature internal to the driver, the Sniffer API presents the rings (or queues) as a first-order interface to the user. Instead of internally maintaining multiple rings as part of the driver, the Sniffer API allows users to directly create rings and retain better control over how rings are allocated on a typical multicore system. The flexibility over ring allocation is balanced with a receive function with much stricter semantics, but only to pursue these goals:

- Allow Sniffer implementation to partition incoming network traffic through receive rings while ensuring that packets maintain order in TCP/UDP flows. This virtualizes the capture interface and allows parallel processing of incoming packet data.
- Give users "zero-copy" access to captured data packets such that users can operate directly on the packet data that was provided by the NIC.
- Minimize the use of locking and other software overheads in the critical path to achieve high packet rates (both on each ring and as an aggregate).

### 5.1.3.1 Multiple Receive Rings

The number of requested receive rings is part of the function call to open a device for packet capture. While users retain control over how many rings are to be allocated by the API, the API best matches processing models that allocate one ring per core, where each core will both capture packets from Sniffer and process/analyze them in place. In-place processing refers to an effort to contain the packet capture and analysis within a CPU core's memory domain, which includes all levels of caches and RAM closest to the core. Since today's multicore platforms are increasingly complex in the non-uniformity of memory access (NUMA), it is best to minimize the amount of memory accesses across memory domains and to maximize the amount of work that can be done by one core within its own memory domain. As such, Sniffer assumes that for best performance, users recognize the importance of opening rings and maintain strong ties between each ring's single consuming thread and the core where it is scheduled.

#### 5.1.3.1.1 Receive Side Scaling Modes

When multiple rings are used, users can specify how packets should be partitioned based on a set of **RSS** flags. These flags instruct how Sniffer should hash incoming packets such that per-flow affinity is maintained in the same ring if so desired. This effectively allows each ring to virtualize the underlying network interface by ensuring that packets within the same flow are deterministically delivered to the same receive ring. Users should not need to reconstruct the ordering of flows even though incoming data is being split across many receive rings.

### 5.1.3.2 Zero-Copy receives

When opening a device with one or multiple rings, users can also specify combined amount of memory that can be consumed by all rings to store packet data. If left unspecified, the implementation will default to choosing a relatively conservative amount of memory, assuming that consumers can process incoming packets at line rate. Resizing the ring should be considered only to address specific buffering concerns and when multiple rings are not possible. It is possible that larger rings are necessary to mitigate the non-real-time behavior of some of the supported operating systems. Larger data rings will only serve as a temporary relief for users that cannot consume incoming data at line rates. If the application does not return packets to the data ring as fast as the packets are coming in, the receive ring will eventually overflow and the packets will be dropped.

### 5.1.3.3 API Software Overhead

The API tries to minimize software overhead in as many areas that can be addressed directly by the library. This includes duplicating some internal data structures to prevent false sharing between multiple ring consumers. More importantly, however, is the explicit decision to **not** provide any locking for the main receive function ([snf\\_ring\\_rcv](#)). Even light forms of locking can impact processing rates when incoming data rates are roughly 15 Mpps (or every 70 nanoseconds). Instead, the API promotes the use of multiple rings to improve concurrency in packet processing and leave more breathing room for packet-per-packet analysis on each consuming core.

### 5.1.3.4 Thread-Safety

Users should consult the API reference for information on thread safety as each function extensively documents how in can be used in threaded environments. While most API functions provide strong thread safety guarantees, the main receive function ([snf\\_ring\\_rcv](#)) is specifically not thread-safe for software overhead reasons explained in [API Software Overhead](#). The expectation is that multiple threads would each open their own receive rings and independently consume packets on their own rings. Specifically, the Sniffer implementation assumes that with each successive call

to `snf_ring_rcv` within the context of a ring always means that the previous packet was consumed by the previous receive on that ring. This should not come as a surprise for users that have been exposed to the serialization present in `libpcap`-type interfaces. However, this may not necessarily match up well with users that expected to separate their computing resources (i.e. cores) between capture and analysis. These users can always resort to using a single receive ring.

## 5.1.4 Implementation details

### 5.1.4.1 Timestamps

Packet timestamps are available from each packet for packet arrival as seen from the NIC. The 64-bit nanosecond since EPOCH timestamp is returned to the user in the receive call. It should be straightforward for the user to convert nanoseconds to a **struct timeval** if so needed. There are 3 timestamping modes: Host, Timesource (Hardware) and Arista. In Host timestamping mode, the timestamp returned to the user in the receive call is normalized in host nanoseconds `ctime` and Sniffer internally ensures that both clocks remain synchronized at a regular interval. The frequency of the NIC's clock is 2MHz plus/minus 100ppm. Timesource (Hardware) mode is available with 10G-PCI-E2-8C2-2S-SYNC adapters when connected to an IRIG-B time source. Arista mode is available when the NIC port is connected to an Arista 71xx series switch port with "FCS append mode" timestamping enabled and the mode is also enabled in Sniffer10G with the `MYRI_ARISTA_ENABLE_TIMESTAMPING=1` environment variable.

### 5.1.4.2 Library Memory Consumption

In order to efficiently use host memory, data rings are allocated when the device is opened as a pool of 4KB pages. At open time, users can specify the number of data rings that are allocated as well as the amount of `data_ring_size` that can be allocated for **all** rings. Internally, the Sniffer library requires an additional 1/32nd of the `data_ring_size` for each ring and some additional inter-ring synchronization memory. This can be summarized as at most 10% of the `data_ring_size` is allocated internally by Sniffer.

### 5.1.4.3 Ethernet and Sniffer Driver Modes

At any time, irrespective of whether the underlying device is enabled for Sniffer, users can configure a Sniffer-capable device as a regular ethernet device (i.e. typically via `ifconfig`). Unless the device is opened for capture, both send and receive functionality work as expected from any Ethernet driver. When the device is opened for capture, packets usually destined to the Ethernet driver are delivered to a Sniffer receive ring instead. However, users can still rely on the Ethernet's RAW sockets interface to send packets (a sample raw socket send enabled with Sniffer receive calls is provided as a test).

## 5.1.5 Define Documentation

### 5.1.5.1 `#define SNF_VERSION_API 8`

SNF API version number (16 bits) Least significant byte increases for minor backwards compatible changes in the API. Most significant byte increases for incompatible changes in the API

0x0008: Add link speed support.

0x0007: Add Multiple Application support and `snf_set_app_id()` function.

0x0006: Internal driver/library API changes.

0x0005: Internal driver/library API changes.

0x0004: Add more injection support and aggregate port opens

0x0003: Add injection support and 3 send counters in statistics.

0x0002: Add `nic_bytes_recv` counter to stats to help users calculate the amount of bandwidth that is actually going through the NIC port.

## 5.1.6 Typedef Documentation

### 5.1.6.1 `typedef struct snf_handle* snf_handle_t`

opaque snf device handle

Opaque snf handle, allocated at [snf\\_open](#) time when a device can be successfully opened

### 5.1.6.2 `typedef struct snf_ring* snf_ring_t`

opaque snf ring handle

Opaque snf ring handle, allocated at [snf\\_ring\\_open](#) time when a ring can be successfully opened. The ring itself can represent a single or an aggregate of physical rings.

## 5.1.7 Enumeration Type Documentation

### 5.1.7.1 `enum snf_link_state`

Link state enumeration, returned by [snf\\_get\\_link\\_state](#)

### 5.1.7.2 `enum snf_timesource_state`

Timesource state (for -SYNC NICs), returned by [snf\\_get\\_timesource\\_state](#)

#### Enumerator:

***SNF\_TIMESOURCE\_LOCAL*** Local timesource (no external). Returned if there is no available external timesource or if its use was explicitly disabled.

***SNF\_TIMESOURCE\_EXT\_UNSYNCED*** External Timesource: not synchronized (yet)

***SNF\_TIMESOURCE\_EXT\_SYNCED*** External Timesource: synchronized

***SNF\_TIMESOURCE\_EXT\_FAILED*** External Timesource: NIC failure to connect to source

***SNF\_TIMESOURCE\_ARISTA\_ACTIVE*** Arista switch is sending ptp timestamps

***SNF\_TIMESOURCE\_PPS*** PPS is being used for time

## 5.1.8 Function Documentation

### 5.1.8.1 `int snf_close ( snf_handle_t devhandle )`

Close port.

This function can be closed once all opened rings (if any) are closed through [snf\\_ring\\_close](#). Once a port is determined to be closable, it is implicitly called as if a call had been previously made to [snf\\_stop](#).

#### Return values

<i>0</i>	Successful.
<i>EBUSY</i>	Some rings are still opened and the port cannot be closed (yet).

#### Postcondition

If successful, all resources allocated at open time are unallocated and the device switches from Sniffer mode to Ethernet mode such that the Ethernet driver resumes receiving packets.

#### 5.1.8.2 void snf\_freeifaddrs ( struct snf\_ifaddrs \* ifaddrs )

Free the list of library allocated memory for [snf\\_getifaddrs](#)

#### Parameters

<i>ifaddrs</i>	Pointer to ifaddrs allocated via <a href="#">snf_getifaddrs</a>
----------------	---

#### 5.1.8.3 int snf\_get\_link\_speed ( snf\_handle\_t devhandle, uint64\_t \* speed )

Get link speed on opened handle

#### Parameters

<i>devhandle</i>	Device handle
<i>speed</i>	Returns speed in bits-per-second for the link

#### Remarks

The cost of retrieving the link speed requires a function call that reads information kept in kernel host memory (i.e. no PCI bus reads).

#### 5.1.8.4 int snf\_get\_link\_state ( snf\_handle\_t devhandle, enum snf\_link\_state \* state )

Get link status on opened handle

#### Parameters

<i>devhandle</i>	Device handle
<i>state</i>	Returns one of SNF_LINK_DOWN or SNF_LINK_UP

#### Remarks

The cost of retrieving the link state requires a function call that reads state kept in kernel host memory (i.e. no PCI bus reads).

#### 5.1.8.5 int snf\_get\_timesource\_state ( snf\_handle\_t devhandle, enum snf\_timesource\_state \* state )

Get Timesource information from opened handle

##### Parameters

<i>devhandle</i>	Device handle
<i>state</i>	Returns one of <a href="#">snf_timesource_state</a>

##### Remarks

The cost of retrieving the timesource state requires a function call that reads state kept in kernel host memory (i.e. no PCI bus reads).

#### 5.1.8.6 int snf\_getifaddrs ( struct snf\_ifaddrs \*\* ifaddrs\_o )

Get a list of Sniffer-capable ethernet devices.

##### Parameters

<i>ifaddrs_o</i>	Library-allocated list of Sniffer-capable devices
------------------	---

##### Remarks

Much like [getifaddrs](#), the user can traverse the list until [snf\\_ifa\\_next](#) is NULL. The interface will show up if the ethernet driver sees the device but the interface does not have to be brought up with an IP address (i.e. no need to 'ifconfig up').

##### Postcondition

User should call [snf\\_freeifaddrs](#) to free the memory that was allocated by the library.

#### 5.1.8.7 int snf\_getportmask\_linkup ( uint32\_t \* mask\_o, int \* cnt\_o )

Get a mask of all Sniffer-capable ports that have their link state set to UP

The least significant bit represents port 0.

Similar to [snf\\_getportmask\\_valid](#) except that only ports with an active link are set in the mask.

##### Parameters

<i>mask_o</i>	bitmask set at output
<i>cnt_o</i>	Number of bits set in bitmask

##### Return values

0	Successful.
ENODEV	Error obtaining port information

### 5.1.8.8 int snf\_getportmask\_valid ( uint32\_t \* mask\_o, int \* cnt\_o )

Get a mask of all Sniffer-capable ports.

The least significant bit represents port 0.

#### Parameters

<i>mask_o</i>	bitmask set at output
<i>cnt_o</i>	Number of bits set in bitmask

#### Return values

0	Successful.
ENODEV	Error obtaining port information

### 5.1.8.9 int snf\_init ( uint16\_t api\_version )

Initialize Sniffer Library with api\_version == [SNF\\_VERSION\\_API](#).

Initializes the sniffer library.

#### Parameters

<i>api_version</i>	Must always be <a href="#">SNF_VERSION_API</a>
--------------------	--

#### Remarks

This must be called before any other call to the sniffer library.

The library may be safely initialized multiple times, although the api\_version should be the same SNF\_VERSION\_API each time.

### 5.1.8.10 int snf\_open ( uint32\_t portnum, int num\_rings, const struct snf\_rss\_params \* rss\_params, int64\_t dataring\_sz, int flags, snf\_handle\_t \* devhandle )

Open device for single or multi-ring operation.

Opens a port for sniffing and allocates a device handle.

#### Parameters

<i>portnum</i>	Port numbers can be interpreted as integers for a specific port number or as a mask when SNF_F_AGGREGATE_PORTMASK is specified in flags. Port information can be obtained through <a href="#">snf_getifaddrs</a> and active/valid masks are available with <a href="#">snf_getportmask_valid</a> and <a href="#">snf_getportmask_linkup</a> . As a special case, if portnum -1 is passed, the library will internally open a portmask as if <a href="#">snf_getportmask_valid</a> was called.
<i>num_rings</i>	Number of rings to allocate for receive-side scaling feature, which determines how many different threads can open their own ring via <a href="#">snf_ring_open()</a> . If set to 0 or less than zero, default value is used unless SNF_NUM_RINGS is set in the environment.

<i>rss_params</i>	Points to a user-initialized structure that selects the RSS mechanism to apply to each incoming packet. This parameter is only meaningful if there are more than 1 rings to be opened. By default, if users pass a NULL value, the implementation will select its own mechanism to divide incoming packets across rings. RSS parameters are documented in <a href="#">Receive-Side Scaling (RSS)</a> .
<i>dataring_sz</i>	Represents the total amount of memory to be used to store incoming packet data for <i>all</i> rings to be opened. If the value is set to 0 or less than 0, the library tries to choose a sensible default unless SNF_DATARING_SIZE is set in the environment. The value can be specified in megabytes (if it is less than 1048576) or is otherwise considered to be in bytes. In either case, the library may slightly adjust the user's request to satisfy alignment requirements (typically 2MB boundaries).
<i>flags</i>	A mask of flags documented in <a href="#">Open flags for process-sharing, port aggregation and packet duplication</a> .
<i>devhandle</i>	Device handle allocated if the call is successful

### Return values

<i>0</i>	Successful. the port is opened and a value devhandle is allocated (see remarks)
<i>EBUSY</i>	Device is already opened
<i>EINVAL</i>	Invalid argument passed, most probably num_rings (if not, check syslog)
<i>E2BIG</i>	Driver could not allocate requested dataring_sz (check syslog)
<i>ENOMEM</i>	Either library or driver did not have enough memory to allocate handle descriptors (but not data ring).
<i>ENODEV</i>	Device portnum can't be opened

### Postcondition

If successful, the NIC switches from Ethernet mode to Capture mode and the Ethernet driver stops receiving packets.

If successful, a call to [snf\\_start](#) is required to the Sniffer-mode NIC to deliver packets to the host, and this call must occur after at least one ring is opened ([snf\\_ring\\_open](#)).

#### 5.1.8.11 int snf\_open\_defaults ( uint32\_t portnum, snf\_handle\_t \* devhandle )

Open device for single or multi-ring operation.

Opens a port for sniffing and allocates a device handle using system defaults.

This function is a simplified version of [snf\\_open](#) and ensures that the resulting device is opened according to system defaults. Since the number of rings and flags can be set by module parameters, some installations may prefer to control device-level parameters in a system-wide configuration and keep the library calls simple.

This call is equivalent to

```
snf_open(portnum, 0, NULL, 0, -1, devhandle);
```

### Parameters

<i>portnum</i>	Ports are numbered from 0 to N-1 where 'N' is the number of Myricom ports available on the system. <a href="#">snf_getifaddrs()</a> may be a useful utility to retrieve the port number by interface name or mac address if there are multiple
<i>devhandle</i>	Device handle allocated if the call is successful



## See also

[snf\\_open](#)

### 5.1.8.12 int snf\_ring\_close ( snf\_ring\_t ringh )

Close a ring

This function is used to inform the underlying device that no further calls to [snf\\_ring\\_recv](#) will be made. If the device is not subsequently closed ([snf\\_close](#)), all packets that would have been delivered to this ring are dropped. Also, by calling this function, users confirm that all packet processing for packets obtained on this ring via [snf\\_ring\\_recv](#) is complete.

#### Parameters

<i>ringh</i>	Ring handle
--------------	-------------

#### Return values

0	Successful.
---	-------------

#### Postcondition

The user has processed the last packet obtained with [snf\\_ring\\_recv](#) and the device can safely be closed via [snf\\_close](#) if all other rings are also closed.

### 5.1.8.13 int snf\_ring\_getstats ( snf\_ring\_t ringh, struct snf\_ring\_stats \* stats )

Get statistics from a receive ring.

#### Parameters

<i>ringh</i>	Ring handle
<i>stats</i>	User-provided pointer to a statistics structure <a href="#">snf_ring_stats</a> , filled in by the library.

#### Remarks

This call is provided as a convenience and should not be relied on for time-critical applications or for high levels of accuracy. Statistics are only updated by the NIC periodically.

#### Warning

Administrative clearing of NIC counters while a Sniffer-based application is running may cause some of the counters to be incorrect.

### 5.1.8.14 int snf\_ring\_open ( snf\_handle\_t devhandle, snf\_ring\_t \* ringh )

Opens the next available ring

### Parameters

<i>devhandle</i>	Device handle, obtained from a successful call to <a href="#">snf_open</a>
<i>ringh</i>	Ring handle allocated if the call is successful.

### Return values

<i>0</i>	Successful. The ring is opened and ringh contains the ring handle.
<i>EBUSY</i>	Too many rings already opened

### Remarks

This function will consider the value of the SNF\_RING\_ID environment variable. For more control over ring allocation, consider using [snf\\_ring\\_open\\_id](#) instead.

### Postcondition

If successful, a call to [snf\\_start](#) is required to the Sniffer-mode NIC to deliver packets to the host.

#### 5.1.8.15 int [snf\\_ring\\_open\\_id](#) ( [snf\\_handle\\_t](#) *devhandle*, int *ring\_id*, [snf\\_ring\\_t](#) \* *ringh* )

Opens a ring from an opened port.

### Parameters

<i>devhandle</i>	Device handle, obtained from a successful call to <a href="#">snf_open</a>
<i>ring_id</i>	Ring number to open, from 0 to <b>num_rings</b> - 1. If the value is -1, this function behaves as if <a href="#">snf_ring_open</a> was called.
<i>ringh</i>	Ring handle allocated if the call is successful.

### Return values

<i>0</i>	Successful. The ring is opened and ringh contains the ring handle.
<i>EBUSY</i>	If ring_id == -1, Too many rings already opened. If ring_id >= 0, that ring is already opened.

### Remarks

Unlike [snf\\_ring\\_open](#) this function ignores the environment variable SNF\_RING\_ID since the expectation is that users want to directly control ring allocation (unlike through libpcap).

### Postcondition

If successful, a call to [snf\\_start](#) is required to the Sniffer-mode NIC to deliver packets to the host.

#### 5.1.8.16 int [snf\\_ring\\_portinfo](#) ( [snf\\_ring\\_t](#) *ring*, [struct snf\\_ring\\_portinfo](#) \* *portinfo* )

Returns information for the ring. For aggregated rings, returns information for each of the physical rings. It is up to the user to make sure they have allocated enough memory to hold the information for all the physical rings in an aggregated ring.

### Parameters

<i>ring</i>	Ring handle (from <a href="#">snf_ring_open</a> )
<i>portinfo</i>	Pointer to memory allocated by the user that will be filled in with the information.

#### 5.1.8.17 int snf\_ring\_recv ( snf\_ring\_t ringh, int timeout\_ms, struct snf\_recv\_req \* recv\_req )

Receive next packet from a receive ring.

This function is used to return the next available packet in a receive ring. The function can block indefinitely, for a specific timeout or be used as a non-blocking call with a timeout of 0.

### Parameters

<i>ringh</i>	Ring handle (from <a href="#">snf_ring_open</a> )
<i>timeout_ms</i>	Receive timeout to control how the function blocks for the next packet. If the value is less than 0, the function can block indefinitely. If the value is 0, the function is guaranteed to never enter a blocking state and returns EAGAIN unless there is a packet waiting. If the value is greater than 0, the caller indicates a desired wait time in milliseconds. With a non-zero wait time, the function only blocks if there are no outstanding packets. If the timeout expires before a packet can be received, the function returns EAGAIN (and <b>not</b> ETIMEDOUT). In all cases, users should expect that the function may return EINTR as the result of signal delivery.
<i>recv_req</i>	Receive Packet structure, only updated when a the function returns 0 for a successful packet receive ( <a href="#">snf_recv_req</a> )

### Return values

<i>0</i>	Successful packet delivery, recv_req is updated with packet information.
<i>EINTR</i>	The call was interrupted by a signal handler
<i>EAGAIN</i>	No packets available (only when timeout is >= 0).

### Remarks

The packet returned always points directly into the receive ring where the NIC has DMAed the packet (there are no copies). As such, the user obtains a pointer to library/driver allocated memory. Users can modify the contents of the packets but should remain within the boundaries of **pkt\_addr** and **length**.

Upon calling the function, the library assumes that the user is done processing the previous packet. The same assumption is made when the ring is closed ([snf\\_ring\\_close](#)).

#### 5.1.8.18 int snf\_ring\_recv\_many ( snf\_ring\_t ring, int timeout\_ms, struct snf\_recv\_req \* req\_vector, int nreq\_in, int \* nreq\_out, struct snf\_ring\_qinfo \* qinfo )

Receive and borrow many packets at once.

This function allows callers to receive one or more packets per call. Contrary to [snf\\_ring\\_recv](#), this function assumes that callers will split the functionality to receive packets (or borrow them) and the functionality to return packets through [snf\\_ring\\_return\\_many](#).

### Parameters

<i>ring</i>	Ring handle (from <a href="#">snf_ring_open</a> )
-------------	---

<i>timeout_ms</i>	Receive timeout to control how the function blocks for the next packet. See complete documentation in <a href="#">snf_ring_rcv</a> .
<i>req_vector</i>	Vector of receive packet structures provided by the user and only updated when a packet is received.
<i>nreq_in</i>	Number of receive packet structures provided in <b>req_vector</b> . No more than <b>nreq_in</b> packets can be received.
<i>nreq_out</i>	Output value for the number of packets actually received and updated in <b>req_vector</b>
<i>qinfo</i>	If non-NULL, the qinfo structure is updated before the function returns 0 or EAGAIN (the function is not updated for other error conditions).  // See <a href="#">snf_ring_return_many</a> documentation for examples

#### 5.1.8.19 int snf\_ring\_return\_many ( snf\_ring\_t ring, uint32\_t data\_qlen, struct snf\_ring\_qinfo \* qinfo )

Return packet space to receive ring.

Under the borrow-many-return-many receive model, it is up to the user to return space in the receive ring. The user achieves this by accumulating packet lengths from the **length\_data** parameter from each packet received and returning the space through this function call.

#### Parameters

<i>ring</i>	Ring handle (from <a href="#">snf_ring_open</a> )
<i>data_qlen</i>	Amount of data returned by previously consumed packets. As a special case, if the value -1 is provided, all data previously borrowed through <a href="#">snf_ring_rcv_many</a> will be returned.
<i>qinfo</i>	If non-NULL, the qinfo structure is updated before the function returns.

```
// Example that shows how the borrow-many-return-many receive model
// works. Since the underlying data is a ring, we don't return a
// packet count, we return data space. The library function call
// overhead can be amortized for high packet rates.
void
rcv_dispatch(snf_ring_t ringh, void (*handler)(void *pkt, uint32_t length)
{
    struct snf_rcv_req reqs[32];
    int rc, nreqs;
    int i, wait_msec = 1000;
    extern int do_exit;
    uint32_t return_length = 0;

    while (1) {
        // Wait up to 1 second for at least one packet to arrive with a
        // maximum of 32 packets within a single call
        rc = snf_ring_rcv_many(wrk->hring, wait_msec, reqs, 32, &nreqs, NULL);
        if (rc == 0) {
            for (i = 0; i < nreqs; i++) {
                // We handle each packet separately and accumulate the amount
                // of data each packet consumes in the ring. Because of
                // alignments length_data is somewhat larger than the packet
                // length
                handler(reqs[i].pkt_addr, reqs[i].length);
                return_length += reqs[i].length_data;
            }
            // We return the data in a single call. We could have gathered
```

```

        // some queue information but not this time around (qinfo set
        // to NULL)
        rc = snf_ring_return_many(wrk->hring, return_length, NULL);
        assert(rc == 0);
        return_length = 0;
    }
    else if (rc == EINTR)
        if (do_exit)
            break;
    }
}

```

#### 5.1.8.20 int snf\_set\_app\_id ( int32\_t id )

Set the application ID.

Sets the application ID.

The user may set the application ID after the call to [snf\\_init](#), but before [snf\\_open](#). When the application ID is set, Sniffer duplicates receive packets to multiple applications. Each application must have a unique ID. Then, each application may utilize a different number of rings. The application can be a process with multiple rings and threads. In this case all rings have the same ID. Or, multiple processes may share the same application ID.

The user may store the application ID in the environment variable SNF\_APP\_ID, instead of calling this function. Both actions have the same effect. SNF\_APP\_ID overrides the ID set via snf\_set\_app\_id.

The user may not run a mix of processes with valid application IDs (not -1) and processes with no IDs (-1). Either all processes have valid IDs or none of them do.

##### Parameters

<i>id</i>	A 32-bit signed integer representing the application ID. A valid ID is any value except -1. -1 is reserved and represents "no ID".
-----------	--

##### Return values

<i>0</i>	Successful
<i>EINVAL</i>	snf_init has not been called or id is -1.

#### 5.1.8.21 int snf\_start ( snf\_handle\_t devhandle )

Start packet capture on a port. Packet capture is only started if it is currently stopped or has not yet started for the first time.

##### Parameters

<i>devhandle</i>	Device handle
------------------	---------------

##### Remarks

It is safe to restart packet capture via [snf\\_start](#) and [snf\\_stop](#).  
This call must be called before any packet can be received.

**5.1.8.22 int snf\_stop ( snf\_handle\_t devhandle )**

Stop packet capture on a port. This function should be used carefully in multi-process mode as a single stop command stops packet capture on all rings. It is usually best to simply [snf\\_ring\\_close](#) a ring to stop capture on a ring.

**Parameters**

<i>devhandle</i>	Device handle
------------------	---------------

**Remarks**

Stop instructs the NIC to drop all packets until the next [snf\\_start\(\)](#) or until the port is closed. The NIC only resumes delivering packets when the port is closed, not when traffic is stopped.

## 5.2 Receive-Side Scaling (RSS)

### Data Structures

- struct [snf\\_rss\\_mode\\_function](#)
- struct [snf\\_rss\\_params](#)

### Defines

- #define [SNF\\_RSS\\_IPV4](#) [SNF\\_RSS\\_IP](#)

### Enumerations

- enum [snf\\_rss\\_params\\_mode](#) { [SNF\\_RSS\\_FLAGS](#) = 0, [SNF\\_RSS\\_FUNCTION](#) = 1 }
- enum [snf\\_rss\\_mode\\_flags](#) {  
    [SNF\\_RSS\\_IP](#) = 0x01, [SNF\\_RSS\\_SRC\\_PORT](#) = 0x10, [SNF\\_RSS\\_DST\\_PORT](#) = 0x20, [SNF\\_RSS\\_GTP](#) =  
    0x40,  
    [SNF\\_RSS\\_GRE](#) = 0x80 }

#### 5.2.1 Detailed Description

These options can be passed as parameters to [snf\\_open](#) when RSS is used.

#### 5.2.2 Define Documentation

##### 5.2.2.1 #define [SNF\\_RSS\\_IPV4](#) [SNF\\_RSS\\_IP](#)

Alias for [SNF\\_RSS\\_IP](#) since IPv4 and IPv6 are always both enabled

#### 5.2.3 Enumeration Type Documentation

##### 5.2.3.1 enum [snf\\_rss\\_mode\\_flags](#)

RSS parameters for [SNF\\_RSS\\_FLAGS](#), flags that can be specified to let the implementation know which fields are significant when generating the hash. By default, RSS is computed on IPv4/IPv6 addresses and source/destination ports when the protocol is TCP or UDP or SCTP, the equivalent of which would be

```
struct sniff_rss_params rssp;  
rssp.mode = SNF_RSS_FLAGS;  
rssp.params.rss_flags = SNF_RSS_IP | SNF_RSS_SRC_PORT | SNF_RSS_DST_PORT;  
  
snf_handle_t hsnf;  
int rc = sniff_open(0, 0, &rssp, 0, -1, &hsnf);  
if (rc == 0)  
    printf("RSS will be applied to IP addresses and TCP/UDP ports if applicable"  
        );
```

**Enumerator:**

*SNF\_RSS\_IP* Include IP (v4 or v6) SRC/DST addr in hash

*SNF\_RSS\_SRC\_PORT* Include TCP/UDP/SCTP SRC port in hash

*SNF\_RSS\_DST\_PORT* Include TCP/UDP/SCTP DST port in hash

*SNF\_RSS\_GTP* Include GTP TEID in hash

*SNF\_RSS\_GRE* Include GRE contents in hash

**5.2.3.2 enum snf\_rss\_params\_mode**

RSS select mode

**Enumerator:**

*SNF\_RSS\_FLAGS* Apply RSS using specified flags

*SNF\_RSS\_FUNCTION* Apply RSS using user-defined function: Kernel API only



## 5.3 Open flags for process-sharing, port aggregation and packet duplication

### Defines

- `#define SNF_F_PSHARED 0x1`
- `#define SNF_F_AGGREGATE_PORTMASK 0x2`
- `#define SNF_F_RX_DUPLICATE 0x300`

### 5.3.1 Detailed Description

These options are passed to the **flags** parameter in [snf\\_open](#) to enable various receive mechanisms.

### 5.3.2 Define Documentation

#### 5.3.2.1 `#define SNF_F_AGGREGATE_PORTMASK 0x2`

Device can be opened for port aggregation (or merging). When this flag is passed, the **portnum** parameter in [snf\\_open](#) is interpreted as a bitmask where each set bit position represents a port number. The Sniffer library will then attempt to open every portnum with its bit set in order to merge the incoming data to the user from multiple ports. Subsequent calls to [snf\\_ring\\_open](#) return a ring handle that internally opens a ring on all underlying ports.

#### 5.3.2.2 `#define SNF_F_PSHARED 0x1`

Device can be process-sharable. This allows multiple independent processes to share rings on the capturing device. This option can be used to design a custom capture solution but is also used in libpcap when multiple rings are requested. In this scenario, each libpcap device sees a fraction of the traffic if multiple rings are used unless the [SNF\\_F\\_RX\\_DUPLICATE](#) option is used, in which case each libpcap device sees the same incoming packets.

#### 5.3.2.3 `#define SNF_F_RX_DUPLICATE 0x300`

Device can duplicate packets to multiple rings as opposed to applying RSS in order to split incoming packets across rings. Users should be aware that with N rings opened, N times the link bandwidth is necessary to process incoming packets without drops. The duplication happens in the host rather than the NIC, so while only up to 10Gbits of traffic crosses the PCIe, N times that bandwidth is necessary on the host.

When duplication is enabled, RSS options are ignored since every packet is delivered to every ring.

## 5.4 Packet injection

### Data Structures

- struct [snf\\_pkt\\_fragment](#)  
*Fragment for [snf\\_inject\\_send\\_v](#).*
- struct [snf\\_inject\\_stats](#)

### Typedefs

- typedef struct [snf\\_inject\\_handle](#) \* [snf\\_inject\\_t](#)

### Functions

- int [snf\\_inject\\_open](#) (int portnum, int flags, [snf\\_inject\\_t](#) \*handle)  
*Open a port for injection and allocate an injection handle.*
- int [snf\\_get\\_injection\\_speed](#) ([snf\\_inject\\_t](#) devhandle, uint64\_t \*speed)
- int [snf\\_inject\\_send](#) ([snf\\_inject\\_t](#) inj, int timeout\_ms, int flags, const void \*pkt, uint32\_t length)  
*Send a packet and optionally block until send resources are available.*
- int [snf\\_inject\\_sched](#) ([snf\\_inject\\_t](#) inj, int timeout\_ms, int flags, const void \*pkt, uint32\_t length, uint64\_t delay\_ns)  
*Send a packet with hardware delay and optionally block until send resources are available.*
- int [snf\\_inject\\_send\\_v](#) ([snf\\_inject\\_t](#) inj, int timeout\_ms, int flags, struct [snf\\_pkt\\_fragment](#) \*frags\_vec, int nfrags, uint32\_t length\_hint)  
*Send a packet assembled from a vector of fragments and optionally block until send resources are available.*
- int [snf\\_inject\\_sched\\_v](#) ([snf\\_inject\\_t](#) inj, int timeout\_ms, int flags, struct [snf\\_pkt\\_fragment](#) \*frags\_vec, int nfrags, uint32\_t length\_hint, uint64\_t delay\_ns)  
*Send a packet assembled from a vector of fragments at a scheduled point relative to the start of the prior packet and optionally block until send resources are available.*
- int [snf\\_inject\\_close](#) ([snf\\_inject\\_t](#) inj)  
*Close injection handle.*
- int [snf\\_inject\\_getstats](#) ([snf\\_inject\\_t](#) inj, struct [snf\\_inject\\_stats](#) \*stats)  
*Get statistics from an injection handle.*

#### 5.4.1 Detailed Description

SNF Packet injection routines that can be used for independent packet generation or coupled to reinject packets received with [snf\\_ring\\_recv](#).

#### 5.4.2 Typedef Documentation

##### 5.4.2.1 typedef struct [snf\\_inject\\_handle](#)\* [snf\\_inject\\_t](#)

Opaque injection handle, allocated by [snf\\_inject\\_open](#). There are only a limited amount of injection handles per NIC/port.

### 5.4.3 Function Documentation

#### 5.4.3.1 `int snf_get_injection_speed ( snf_inject_t devhandle, uint64_t * speed )`

Get link speed on opened injection handle

##### Parameters

<i>devhandle</i>	Device handle
<i>speed</i>	Returns speed in bits-per-second for the link

##### Remarks

The cost of retrieving the link speed requires a function call that reads information kept in kernel host memory (i.e. no PCI bus reads).

#### 5.4.3.2 `int snf_inject_close ( snf_inject_t inj )`

Close injection handle.

This function closes an injection handle and ensures that all pending sends are sent by the NIC.

##### Parameters

<i>inj</i>	Injection handle
------------	------------------

##### Return values

0	Successful
---	------------

##### Postcondition

Once closed, the injection handle will have ensured that any pending sends have been sent out on the wire. The handle is then made available again for the underlying port's limited amount of handles.

#### 5.4.3.3 `int snf_inject_getstats ( snf_inject_t inj, struct snf_inject_stats * stats )`

Get statistics from an injection handle.

##### Parameters

<i>inj</i>	Injection Handle
<i>stats</i>	User-provided pointer to a statistics structure <a href="#">snf_inject_stats</a> , filled in by the SNF implementation.

##### Remarks

This call is provided as a convenience and should not be relied on for time-critical applications or for high levels of accuracy. Statistics are only updated by the NIC periodically.

## Warning

Administrative clearing of NIC counters while a Sniffer-based application is running may cause some of the counters to be incorrect.

### 5.4.3.4 int snf\_inject\_open ( int portnum, int flags, snf\_inject\_t \* handle )

Open a port for injection and allocate an injection handle.

#### Parameters

<i>portnum</i>	Ports are numbered from 0 to N-1 where 'N' is the number of Myricom ports available on the system. <a href="#">snf_getifaddrs()</a> may be a useful utility to retrieve the port number by interface name or mac address if there are multiple
<i>flags</i>	Flags for injection handle. None are currently defined.
<i>handle</i>	Injection handle allocated if the call is successful.

#### Return values

<i>0</i>	Success. An injection handle is opened and allocated.
<i>EBUSY</i>	Ran out of injection handles for this port
<i>ENOMEM</i>	Ran out of memory to allocate new injection handle

### 5.4.3.5 int snf\_inject\_sched ( snf\_inject\_t inj, int timeout\_ms, int flags, const void \* pkt, uint32\_t length, uint64\_t delay\_ns )

Send a packet with hardware delay and optionally block until send resources are available.

This send function is used for paced packet injection. This function can be used as part of a packet replay program. When the function returns successfully, the packet is guaranteed to be completely buffered by SNF: no references are kept to the input data and the caller is free to safely modify its contents. The SNF implementation delays transmitting the packet according to the *delay\_ns* parameter, relative to the start of the prior packet.

#### Parameters

<i>inj</i>	Injection handle
<i>timeout_ms</i>	Timeout in milliseconds to wait if insufficient send resources are available to inject a new packet. Insufficient resources can be a lack of send descriptors or a full send queue ring. If <i>timeout_ms</i> is 0, the function won't block for send resources and returns EAGAIN.
<i>flags</i>	Flags (currently none).
<i>pkt</i>	Pointer to the packet to be sent. The packet must be a pointer to a complete Ethernet frame (without the trailing CRC) and start with a valid Ethernet header. The hardware will append 4-CRC bytes at the end of the packet. The maximum valid packet size is 9000 bytes and is enforced by the library. The minimum valid packet size is 60 bytes, although any packet smaller than 60 bytes will be accepted by the library and padded by the hardware.
<i>length</i>	The length of the packet, excluding the trailing 4 CRC bytes.
<i>delay_ns</i>	The minimum delay between the start of the prior packet and the start of this packet. Packets with a delay less than the time to send the prior packet are send immediately. It is recommended to use 0 as the delta on the first packet sent.

### Return values

<i>0</i>	Successful. The packet is buffered by SNF.
<i>EAGAIN</i>	Insufficient resources to send packet. If <i>timeout_ms</i> is non-zero, the caller will have blocked at least that many milliseconds before resources could become available.
<i>EINVAL</i>	Packet length is larger than 9000 bytes.
<i>ENOTSUP</i>	The hardware does not support injection pacing.

### Postcondition

If successful, the packet is completely buffered for sending by SNF. The implementation guarantees that it will eventually send the packet out, as scheduled, without requiring further calls into SNF.

**5.4.3.6** `int snf_inject_sched_v ( snf_inject_t inj, int timeout_ms, int flags, struct snf_pkt_fragment * frags_vec, int nfrags, uint32_t length_hint, uint64_t delay_ns )`

Send a packet assembled from a vector of fragments at a scheduled point relative to the start of the prior packet and optionally block until send resources are available.

This send function follows the same semantics as [snf\\_inject\\_send](#) except that the packet to be injected can be assembled from multiple fragments (or buffers).

### Parameters

<i>inj</i>	Injection handle
<i>timeout_ms</i>	Timeout in milliseconds to wait if insufficient send resources are available to inject a new packet. Insufficient resources can be a lack of send descriptors or a full send queue ring. If <i>timeout_ms</i> is 0, the function won't block for send resources and returns <i>EAGAIN</i> .
<i>flags</i>	Flags (currently none).
<i>frags_vec</i>	Pointer to a vector of 1 or more buffers/fragments that can be used to compose a complete Ethernet frame (not including the trailing CRC header). The first fragment must point to a valid Ethernet header and the hardware will append its own (valid 4-byte CRC) at the end of the last buffer/fragment passed in the <i>frags_vec</i> . When all the fragments are added up, the maximum valid packet size is 9000 bytes and is enforced by the library. The minimum valid packet size is 60 bytes, although any packet smaller than 60 bytes will be accepted by the library and padded by the hardware.
<i>nfrags</i>	Number of elements in the io vector
<i>length_hint</i>	If non-zero, the amount is expected to be the sum of all the lengths passed in the io vector. This parameters can help the library account for space when injecting packets.
<i>delay_ns</i>	The minimum delay between the start of the prior packet and the start of this packet. Packets with a delay less than the time to send the prior packet are send immediately. It is recommended to use 0 as the delta on the first packet sent.

### Return values

<i>0</i>	Successful. The packet is buffered by SNF.
<i>EAGAIN</i>	Insufficient resources to send packet. If <i>timeout_ms</i> is non-zero, the caller will have blocked at least that many milliseconds before resources could become available.
<i>EINVAL</i>	Packet length is larger than 9000 bytes.
<i>ENOTSUP</i>	The hardware does not support injection pacing.

## Postcondition

If successful, the packet is completely buffered for send by SNF. The implementation guarantees that it will eventually send the packet out in a timely fashion without requiring further calls into SNF.

```
// Example that takes an existing packet and prepends the existing
// ethernet type with a vlan header.
//
int
send_prepend_vlan_tag(uint16_t vtag, void *pkt, uint32_t len)
{
    uint32_t vlanhdr = htonl(0x8100 << 16 | vtag);
    struct snf_pkt_fragment vec[3];

    // We assume that the input 'pkt' does not already contain a vlan tag
    // and that the pkt is not terminated with a CRC. The hardware will
    // add the CRC. We also use no timeout in the send meaning that the
    // send may return EAGAIN if there are insufficient resources to
    // queue the send.

    vec[0].ptr = (void *) pkt;
    vec[0].length = 12; // dest and src mac
    vec[1].ptr = &vlanhdr;
    vec[1].length = sizeof(vlanhdr);
    vec[2].ptr = (void *) ((uint8_t *) pkt + 12);
    vec[2].length = len - 12;
    len += sizeof(vlanhdr);

    // Schedule the packet to be sent 3 us. from the last.
    return snf_inject_sched_v(hinj, 0, 0, vec, 3, len, 3000);
}
```

### 5.4.3.7 int snf\_inject\_send ( snf\_inject\_t inj, int timeout.ms, int flags, const void \* pkt, uint32\_t length )

Send a packet and optionally block until send resources are available.

This send function is optimized for high packet rate injection. While it can be coupled with a receive ring to reinject a packet, it is not strictly necessary. This function can be used as part of a packet generator. When the function returns successfully, the packet is guaranteed to be completely buffered by SNF: no references are kept to the input data and the caller is free to safely modify its contents. A successful return does not, however, guarantee that the packet has been injected into the network. The SNF implementation may choose to hold on to the packet for coalescing in order to improve packet throughput.

## Parameters

<i>inj</i>	Injection handle
<i>timeout_ms</i>	Timeout in milliseconds to wait if insufficient send resources are available to inject a new packet. Insufficient resources can be a lack of send descriptors or a full send queue ring. If <i>timeout_ms</i> is 0, the function won't block for send resources and returns EAGAIN.
<i>flags</i>	Flags (currently none).
<i>pkt</i>	Pointer to the packet to be sent. The packet must be a pointer to a complete Ethernet frame (without the trailing CRC) and start with a valid Ethernet header. The hardware will append 4-CRC bytes at the end of the packet. The maximum valid packet size is 9000 bytes and is enforced by the library. The minimum valid packet size is 60 bytes, although any packet smaller than 60 bytes will be accepted by the library and padded by the hardware.
<i>length</i>	The length of the packet, excluding the trailing 4 CRC bytes.

### Return values

<i>0</i>	Successful. The packet is buffered by SNF.
<i>EAGAIN</i>	Insufficient resources to send packet. If <code>timeout_ms</code> is non-zero, the caller will have blocked at least that many milliseconds before resources could become available.
<i>EINVAL</i>	Packet length is larger than 9000 bytes.

### Postcondition

If successful, the packet is completely buffered for sending by SNF. The implementation guarantees that it will eventually send the packet out in a timely fashion without requiring further calls into SNF.

**5.4.3.8** `int snf_inject_send_v ( snf_inject_t inj, int timeout_ms, int flags, struct snf_pkt_fragment * frags_vec, int nfrags, uint32_t length_hint )`

Send a packet assembled from a vector of fragments and optionally block until send resources are available.

This send function follows the same semantics as [snf\\_inject\\_send](#) except that the packet to be injected can be assembled from multiple fragments (or buffers).

### Parameters

<i>inj</i>	Injection handle
<i>timeout_ms</i>	Timeout in milliseconds to wait if insufficient send resources are available to inject a new packet. Insufficient resources can be a lack of send descriptors or a full send queue ring. If <code>timeout_ms</code> is 0, the function won't block for send resources and returns <i>EAGAIN</i> .
<i>flags</i>	Flags (currently none).
<i>frags_vec</i>	Pointer to a vector of 1 or more buffers/fragments that can be used to compose a complete Ethernet frame (not including the trailing CRC header). The first fragment must point to a valid Ethernet header and the hardware will append its own (valid 4-byte CRC) at the end of the last buffer/fragment passed in the <code>frags_vec</code> . When all the fragments are added up, the maximum valid packet size is 9000 bytes and is enforced by the library. The minimum valid packet size is 60 bytes, although any packet smaller than 60 bytes will be accepted by the library and padded by the hardware.
<i>nfrags</i>	Number of elements in the io vector
<i>length_hint</i>	If non-zero, the amount is expected to be the sum of all the lengths passed in the io vector. This parameters can help the library account for space when injecting packets.

### Return values

<i>0</i>	Successful. The packet is buffered by SNF.
<i>EAGAIN</i>	Insufficient resources to send packet. If <code>timeout_ms</code> is non-zero, the caller will have blocked at least that many milliseconds before resources could become available.
<i>EINVAL</i>	Packet length (in <code>length_hint</code> or the sum of all <code>frags_vec</code> lens) is larger than 9000 bytes.

### Postcondition

If successful, the packet is completely buffered for send by SNF. The implementation guarantees that it will eventually send the packet out in a timely fashion without requiring further calls into SNF.

```
// Example that takes an existing packet and prepends the existing
```

```
// ethernet type with a vlan header.
//
int
send_prepend__vlan_tag(uint16_t vtag, void *pkt, uint32_t len)
{
    uint32_t vlanhdr = htonl(0x8100 << 16 | vtag);
    struct snf_pkt_fragment vec[3];

    // We assume that the input 'pkt' does not already contain a vlan tag
    // and that the pkt is not terminated with a CRC. The hardware will
    // add the CRC. We also use no timeout in the send meaning that the
    // send may return EAGAIN if there are insufficient resources to
    // queue the send.

    vec[0].ptr = (void *) pkt;
    vec[0].length = 12; // dest and src mac
    vec[1].ptr = &vlanhdr;
    vec[1].length = sizeof(vlanhdr);
    vec[2].ptr = (void *) ((uint8_t *)pkt + 12);
    vec[2].length = len - 12;
    len += sizeof(vlanhdr);

    return snf_inject_send_v(hinj, 0, 0, vec, 3, len);
}
```



## 5.5 Packet reflect to netdev (kernel stack)

### Typedefs

- typedef void \* [snf\\_netdev\\_reflect\\_t](#)

### Functions

- int [snf\\_netdev\\_reflect\\_enable](#) ([snf\\_handle\\_t](#) hsnf, [snf\\_netdev\\_reflect\\_t](#) \*handle)  
*Enable a network device for packet reflection.*
- int [snf\\_netdev\\_reflect](#) ([snf\\_netdev\\_reflect\\_t](#) ref\_dev, const void \*pkt, uint32\_t length)  
*Reflect a packet to the network device.*

#### 5.5.1 Detailed Description

Network packets acquired through Sniffer can be reflected back into the kernel path as if the device had initially sent then through to the regular network stack.

While Sniffer users are typically expected to process a significant portion of their packets with less overhead in userspace, this feature is provided as a convenience to allow some packets to be processed back in the kernel. The implementation makes no explicit step to make the kernel-based processing any faster than it is when Sniffer is not being used (in fact, it is probably much slower).

#### 5.5.2 Typedef Documentation

##### 5.5.2.1 typedef void\* [snf\\_netdev\\_reflect\\_t](#)

Opaque handle returned by [snf\\_netdev\\_reflect\\_enable](#) and used to reflect packets onto by [snf\\_netdev\\_reflect](#).

#### 5.5.3 Function Documentation

##### 5.5.3.1 int [snf\\_netdev\\_reflect](#) ( [snf\\_netdev\\_reflect\\_t](#) ref\_dev, const void \* [pkt](#), uint32\_t [length](#) )

Reflect a packet to the network device.

#### Parameters

<i>ref_dev</i>	Reflection handle
<i>pkt</i>	Pointer to the packet to be reflected to the network device. The packet must be a pointer to a complete Ethernet frame (without the trailing CRC) and start with a valid Ethernet header.
<i>length</i>	The length of the packet, excluding the trailing 4 CRC bytes.

#### Return values

0	Successful. The packet is buffered by SNF.
---	--

**Postcondition**

If successful, the packet is completely buffered into the network device receive path.

**5.5.3.2 int snf\_netdev\_reflect\_enable ( snf\_handle\_t *hsnf*, snf\_netdev\_reflect\_t \* *handle* )**

Enable a network device for packet reflection.

**Parameters**

<i>hsnf</i>	handle for network device to reflect onto, obtained by <a href="#">snf_open</a>
<i>handle</i>	Reflection handle.

**Return values**

0	Success. An reflection handle is enabled.
---	---

## **Chapter 6**

# **Namespace Documentation**

### **6.1 snf Namespace Reference**

#### **6.1.1 Detailed Description**

Sniffer

#### **Author**

Myricom, Inc.



## Chapter 7

# Data Structure Documentation

### 7.1 snf\_ifaddrs Struct Reference

#### Data Fields

- struct [snf\\_ifaddrs](#) \* [snf\\_ifa\\_next](#)
- const char \* [snf\\_ifa\\_name](#)
- uint32\_t [snf\\_ifa\\_portnum](#)
- int [snf\\_ifa\\_maxrings](#)
- uint8\_t [snf\\_ifa\\_macaddr](#) [6]
- uint8\_t [pad](#) [2]
- int [snf\\_ifa\\_maxinject](#)
- enum [snf\\_link\\_state](#) [snf\\_ifa\\_link\\_state](#)
- uint64\_t [snf\\_ifa\\_link\\_speed](#)

#### 7.1.1 Detailed Description

Structure to map Interfaces to Sniffer port numbers

#### 7.1.2 Field Documentation

##### 7.1.2.1 uint8\_t snf\_ifaddrs::pad[2]

Internal padding (ignore)

##### 7.1.2.2 uint64\_t snf\_ifaddrs::snf\_ifa\_link\_speed

Link Speed (bps)

**7.1.2.3 enum snf\_link\_state snf\_ifaddr::snf\_ifa\_link\_state**

Underlying port's state (DOWN or UP)

**7.1.2.4 uint8\_t snf\_ifaddr::snf\_ifa\_macaddr[6]**

MAC address

**7.1.2.5 int snf\_ifaddr::snf\_ifa\_maxinject**

Maximum TX injection handles supported

**7.1.2.6 int snf\_ifaddr::snf\_ifa\_maxrings**

Maximum RX rings supported

**7.1.2.7 const char\* snf\_ifaddr::snf\_ifa\_name**

interface name, as in ifconfig

**7.1.2.8 struct snf\_ifaddr\* snf\_ifaddr::snf\_ifa\_next**

next item or NULL if last

**7.1.2.9 uint32\_t snf\_ifaddr::snf\_ifa\_portnum**

snf port number

## **7.2 snf\_inject\_stats Struct Reference**

### **Data Fields**

- uint64\_t [inj\\_pkt\\_send](#)
- uint64\_t [nic\\_pkt\\_send](#)
- uint64\_t [nic\\_bytes\\_send](#)

### **7.2.1 Detailed Description**

Structure to return statistics from an injection handle. The hardware-specific counters (nic\_) apply to all injection handles.

## 7.2.2 Field Documentation

### 7.2.2.1 `uint64_t snf_inject_stats::inj_pkt_send`

Number of packets sent by this injection endpoint

### 7.2.2.2 `uint64_t snf_inject_stats::nic_bytes_send`

Number of raw bytes sent by Hardware Interface (see `nic_bytes_rcv`)

### 7.2.2.3 `uint64_t snf_inject_stats::nic_pkt_send`

Number of total packets sent by Hardware Interface

## 7.3 `snf_pkt_fragment` Struct Reference

Fragment for `snf_inject_send_v`.

### Data Fields

- `const void *` [ptr](#)
- `uint32_t` [length](#)

### 7.3.1 Detailed Description

Fragment for `snf_inject_send_v`.

### 7.3.2 Field Documentation

#### 7.3.2.1 `uint32_t snf_pkt_fragment::length`

Number of bytes

#### 7.3.2.2 `const void* snf_pkt_fragment::ptr`

Packet starting address

## 7.4 `snf_rcv_req` Struct Reference

### Data Fields

- `void *` [pkt\\_addr](#)

- [uint32\\_t length](#)
- [uint64\\_t timestamp](#)
- [uint32\\_t portnum](#)
- [uint32\\_t length\\_data](#)
- [uint32\\_t hw\\_hash](#)

### 7.4.1 Detailed Description

Structure to describe a packet received on a data ring.

### 7.4.2 Field Documentation

#### 7.4.2.1 `uint32_t snf_rcv_req::hw_hash`

Hash calculated by the NIC.

#### 7.4.2.2 `uint32_t snf_rcv_req::length`

Length of packet, does not include Ethernet CRC

#### 7.4.2.3 `uint32_t snf_rcv_req::length_data`

Length of packet, with alignment in receive queue

#### 7.4.2.4 `void* snf_rcv_req::pkt_addr`

Pointer to packet directly in data ring

#### 7.4.2.5 `uint32_t snf_rcv_req::portnum`

Which port number received the packet

#### 7.4.2.6 `uint64_t snf_rcv_req::timestamp`

64-bit timestamp in nanoseconds

## 7.5 `snf_ring_portinfo` Struct Reference

### Data Fields

- [snf\\_ring\\_t ring](#)
- [uintptr\\_t q\\_size](#)
- [uint32\\_t portcnt](#)



- [uint32\\_t portmask](#)
- [uintptr\\_t data\\_addr](#)
- [uintptr\\_t data\\_size](#)

### 7.5.1 Detailed Description

Receive ring information

### 7.5.2 Field Documentation

#### 7.5.2.1 `uintptr_t snf_ring_portinfo::data_addr`

Address of data ring

#### 7.5.2.2 `uintptr_t snf_ring_portinfo::data_size`

Size of the data ring

#### 7.5.2.3 `uint32_t snf_ring_portinfo::portcnt`

How many physical ports deliver to this receive ring

#### 7.5.2.4 `uint32_t snf_ring_portinfo::portmask`

Which ports deliver to this receive ring

#### 7.5.2.5 `uintptr_t snf_ring_portinfo::q_size`

Size of the data queue

#### 7.5.2.6 `snf_ring_t snf_ring_portinfo::ring`

Single ring

## 7.6 `snf_ring_qinfo` Struct Reference

### Data Fields

- [uintptr\\_t q\\_avail](#)
- [uintptr\\_t q\\_borrowed](#)
- [uintptr\\_t q\\_free](#)

### 7.6.1 Detailed Description

Queue consumption information

### 7.6.2 Field Documentation

#### 7.6.2.1 `uintptr_t snf_ring_qinfo::q_avail`

Amount of data available not yet received (approximate)

#### 7.6.2.2 `uintptr_t snf_ring_qinfo::q_borrowed`

Amount of data currently borrowed (exact)

#### 7.6.2.3 `uintptr_t snf_ring_qinfo::q_free`

Amount of free space still available (approximate)

## 7.7 `snf_ring_stats` Struct Reference

### Data Fields

- `uint64_t nic_pkt_rcv`
- `uint64_t nic_pkt_overflow`
- `uint64_t nic_pkt_bad`
- `uint64_t ring_pkt_rcv`
- `uint64_t ring_pkt_overflow`
- `uint64_t nic_bytes_rcv`
- `uint64_t snf_pkt_overflow`
- `uint64_t nic_pkt_dropped`

### 7.7.1 Detailed Description

Structure to return statistics from a ring. The Hardware-specific counters apply to all rings as they are counted before any demultiplexing to a ring is applied.

### 7.7.2 Field Documentation

#### 7.7.2.1 `uint64_t snf_ring_stats::nic_bytes_rcv`

Number of raw bytes received by the Hardware Interface on all rings. Each Ethernet data packet includes 8 bytes of HW header, 4 bytes of CRC and the result is aligned to 16 bytes such that a minimum size 60 byte packet counts for 80 bytes.

**7.7.2.2 uint64\_t snf\_ring\_stats::nic\_pkt\_bad**

Number of Bad CRC/PHY packets seen by Hardware Interface

**7.7.2.3 uint64\_t snf\_ring\_stats::nic\_pkt\_dropped**

Number of packets dropped, reflected in Packets Drop Filter in Counters.

**7.7.2.4 uint64\_t snf\_ring\_stats::nic\_pkt\_overflow**

Number of packets dropped by Hardware Interface

**7.7.2.5 uint64\_t snf\_ring\_stats::nic\_pkt\_recv**

Number of packets received by Hardware Interface

**7.7.2.6 uint64\_t snf\_ring\_stats::ring\_pkt\_overflow**

Number of packets dropped because of insufficient space in receive ring

**7.7.2.7 uint64\_t snf\_ring\_stats::ring\_pkt\_recv**

Number of packets received into the receive ring

**7.7.2.8 uint64\_t snf\_ring\_stats::snf\_pkt\_overflow**

Number of packets dropped because of insufficient space in shared SNF buffering

## **7.8 snf\_rss\_mode\_function Struct Reference**

### **Data Fields**

- `int(* rss\_hash\_fn)(struct snf\_recv\_req *r, void *context, uint32_t *hashval)`
- `void * rss\_context`

### **7.8.1 Detailed Description**

User-defined RSS hashing function parameters. Users that provide their own callbacks can generate their own hash based on the contents of a received packet. **NOTE** This feature is available only in the SNF kernel API

## 7.8.2 Field Documentation

### 7.8.2.1 void\* snf\_rss\_mode\_function::rss\_context

User context that is reflected when the user-provided `rss_hash_fn` is called.

### 7.8.2.2 int(\* snf\_rss\_mode\_function::rss\_hash\_fn)(struct snf\_recv\_req \*r, void \*context, uint32\_t \*hashval)

User-provided hash function. The callback is provided with a valid `snf_recv_req` structure which contains a packet as received by Sniffer. It is up to the user to inspect and parse the packet to produce a unique 32-bit hash. The implementation will map the 32-bit into one of the rings allocated in `snf_open`. The function must return one of three values

- **0** The packet is queued in the ring based on the 32-bit hash value that is provided, which is `hashval % num_rings`.
- **<0** The packet is dropped and accounted as a drop in the ring corresponding to the 32-bit hash value provided by the user.

In the example below, we replace the default hash function with a hash function that sends packets to different rings at every interval of 500 packets. This approach ignores the actual packet contents and the importance of flow affinity, we just want to spread the packet analysis to different rings and threads.

```
#define MAX_RINGS 32
#define PKT_INTERVAL 500
static uint32_t cnt[MAX_RINGS];
static uint32_t cur_ring = 0;

static int
custom_hash(struct snf_recv_req *r, void *context, uint32_t *hashval)
{
    if (++cnt[cur_ring] == PKT_INTERVAL) {
        cnt[cur_ring] = 0;
        if (++cur_ring == MAX_RINGS)
            cur_ring = 0;
    }
    // Return cur_ring as the hash value, since Sniffer will apply a
    // modulo and the corresponding ring will receive the packet when
    // calling snf_recv()
    *hash_val = cur_ring;
    return 0;
}

// At snf_open time, select a custom hash approach.
struct snf_rss_params rssp;
rssp.mode = SNF_RSS_FUNCTION;
rssp.params.rss_function.rss_hash_fn = custom_hash;
rssp.params.rss_function.rss_context = NULL; // Don't need a context

// On port 0, we will open MAX_RINGS rings, use default flags,
// select an 800 MB data ring and chose our custom hashing function.
snf_handle_t hsnf;
int rc = snf_open(0, MAX_RINGS, &rssp, 800, 0, &hsnf);
if (rc) {
    perror("Error in snf_open");
    exit(EXIT_FAILURE);
}
```

## 7.9 snf\_rss\_params Struct Reference

### Data Fields

- enum [snf\\_rss\\_params\\_mode](#) mode
- union {
  - enum [snf\\_rss\\_mode\\_flags](#) rss\_flags
  - struct [snf\\_rss\\_mode\\_function](#) rss\_function
- } [params](#)

### 7.9.1 Detailed Description

When using multiple rings, users can either let Sniffer how to partition the flows of incoming packets or control the hashing using specific RSS modes. The following modes are available.

- **None:** `rss_params` is NULL in [snf\\_open](#). When no RSS mode is explicitly specified, users let the implementation chose an RSS strategy that best matches the revision of the Myri-10G NIC. Unless a specific hashing strategy is required, this approach is best in terms of performance-portability.
- **Flag-based:** `rss_params` sets mode to [SNF\\_RSS\\_FLAGS](#). This mode allows users to functionally specify which parts of a packet are significant in the RSS hashing process. A functional specification leaves enough room for the Sniffer implementation to move part or all of the hash computation between hardware, firmware and software.
- **Function-based** (kernel API only): `rss_params` sets mode to [SNF\\_RSS\\_FUNCTION](#). This mode guarantees the most flexibility for the user but forces the hashing to be serialized in software (note that the current generation NICs do not necessarily take a very large performance hit compared to the two other RSS modes). This approach may be required if the flag-based approach isn't flexible enough. For example, some users may require that flow affinity be maintained according to an encapsulated TCP/IP header. See [snf\\_rss\\_mode\\_function](#) for more details.

### 7.9.2 Field Documentation

#### 7.9.2.1 enum `snf_rss_params_mode` `snf_rss_params::mode`

RSS mode

#### 7.9.2.2 union { ... } `snf_rss_params::params`

RSS parameter settings, according to the mode that is selected

#### 7.9.2.3 enum `snf_rss_mode_flags` `snf_rss_params::rss_flags`

RSS parameters for [SNF\\_RSS\\_FLAGS](#)

**7.9.2.4 struct snf\_rss\_mode\_function snf\_rss\_params::rss\_function**

RSS params for [SNF\\_RSS\\_FUNCTION](#)

# Index

- data\_addr
  - snf\_ring\_portinfo, [43](#)
- data\_size
  - snf\_ring\_portinfo, [43](#)
- hw\_hash
  - snf\_rcv\_req, [42](#)
- inj\_pkt\_send
  - snf\_inject\_stats, [41](#)
- length
  - snf\_pkt\_fragment, [41](#)
  - snf\_rcv\_req, [42](#)
- length\_data
  - snf\_rcv\_req, [42](#)
- mode
  - snf\_rss\_params, [47](#)
- nic\_bytes\_rcv
  - snf\_ring\_stats, [44](#)
- nic\_bytes\_send
  - snf\_inject\_stats, [41](#)
- nic\_pkt\_bad
  - snf\_ring\_stats, [44](#)
- nic\_pkt\_dropped
  - snf\_ring\_stats, [45](#)
- nic\_pkt\_overflow
  - snf\_ring\_stats, [45](#)
- nic\_pkt\_rcv
  - snf\_ring\_stats, [45](#)
- nic\_pkt\_send
  - snf\_inject\_stats, [41](#)
- Open flags for process-sharing, port aggregation and packet duplication, [27](#)
  - SNF\_F\_PSHARED, [27](#)
- Packet injection, [28](#)
  - snf\_get\_injection\_speed, [29](#)
  - snf\_inject\_close, [29](#)
  - snf\_inject\_getstats, [29](#)
  - snf\_inject\_open, [30](#)
  - snf\_inject\_sched, [30](#)
  - snf\_inject\_sched\_v, [31](#)
  - snf\_inject\_send, [32](#)
  - snf\_inject\_send\_v, [33](#)
  - snf\_inject\_t, [28](#)
- Packet reflect to netdev (kernel stack), [35](#)
  - snf\_netdev\_reflect, [35](#)
  - snf\_netdev\_reflect\_enable, [36](#)
  - snf\_netdev\_reflect\_t, [35](#)
- pad
  - snf\_ifaddrs, [39](#)
- params
  - snf\_rss\_params, [47](#)
- pkt\_addr
  - snf\_rcv\_req, [42](#)
- portcnt
  - snf\_ring\_portinfo, [43](#)
- portmask
  - snf\_ring\_portinfo, [43](#)
- portnum
  - snf\_rcv\_req, [42](#)
- ptr
  - snf\_pkt\_fragment, [41](#)
- q\_avail
  - snf\_ring\_qinfo, [44](#)
- q\_borrowed
  - snf\_ring\_qinfo, [44](#)
- q\_free
  - snf\_ring\_qinfo, [44](#)
- q\_size
  - snf\_ring\_portinfo, [43](#)
- Receive-Side Scaling (RSS)
  - SNF\_RSS\_DST\_PORT, [26](#)
  - SNF\_RSS\_FLAGS, [26](#)
  - SNF\_RSS\_FUNCTION, [26](#)
  - SNF\_RSS\_GRE, [26](#)
  - SNF\_RSS\_GTP, [26](#)
  - SNF\_RSS\_IP, [26](#)
  - SNF\_RSS\_SRC\_PORT, [26](#)

- Receive-Side Scaling (RSS), [25](#)
  - SNF\_RSS\_IPV4, [25](#)
  - snf\_rss\_mode\_flags, [25](#)
  - snf\_rss\_params\_mode, [26](#)
- ring
  - snf\_ring\_portinfo, [43](#)
- ring\_pkt\_overflow
  - snf\_ring\_stats, [45](#)
- ring\_pkt\_rcv
  - snf\_ring\_stats, [45](#)
- rss\_context
  - snf\_rss\_mode\_function, [46](#)
- rss\_flags
  - snf\_rss\_params, [47](#)
- rss\_function
  - snf\_rss\_params, [47](#)
- rss\_hash\_fn
  - snf\_rss\_mode\_function, [46](#)
- SNF API Reference
  - SNF\_TIMESOURCE\_ARISTA\_ACTIVE, [14](#)
  - SNF\_TIMESOURCE\_EXT\_FAILED, [14](#)
  - SNF\_TIMESOURCE\_EXT\_SYNCED, [14](#)
  - SNF\_TIMESOURCE\_EXT\_UNSYNCED, [14](#)
  - SNF\_TIMESOURCE\_LOCAL, [14](#)
  - SNF\_TIMESOURCE\_PPS, [14](#)
- SNF\_RSS\_DST\_PORT
  - Receive-Side Scaling (RSS), [26](#)
- SNF\_RSS\_FLAGS
  - Receive-Side Scaling (RSS), [26](#)
- SNF\_RSS\_FUNCTION
  - Receive-Side Scaling (RSS), [26](#)
- SNF\_RSS\_GRE
  - Receive-Side Scaling (RSS), [26](#)
- SNF\_RSS\_GTP
  - Receive-Side Scaling (RSS), [26](#)
- SNF\_RSS\_IP
  - Receive-Side Scaling (RSS), [26](#)
- SNF\_RSS\_SRC\_PORT
  - Receive-Side Scaling (RSS), [26](#)
- SNF\_TIMESOURCE\_ARISTA\_ACTIVE
  - SNF API Reference, [14](#)
- SNF\_TIMESOURCE\_EXT\_FAILED
  - SNF API Reference, [14](#)
- SNF\_TIMESOURCE\_EXT\_SYNCED
  - SNF API Reference, [14](#)
- SNF\_TIMESOURCE\_EXT\_UNSYNCED
  - SNF API Reference, [14](#)
- SNF\_TIMESOURCE\_LOCAL
  - SNF API Reference, [14](#)
- SNF\_TIMESOURCE\_PPS
  - SNF API Reference, [14](#)
- SNF API Reference, [14](#)
- SNF API Reference, [9](#)
- SNF\_VERSION\_API, [13](#)
  - snf\_close, [14](#)
  - snf\_freeifaddrs, [15](#)
  - snf\_get\_link\_speed, [15](#)
  - snf\_get\_link\_state, [15](#)
  - snf\_get\_timesource\_state, [15](#)
  - snf\_getifaddrs, [16](#)
  - snf\_getportmask\_linkup, [16](#)
  - snf\_getportmask\_valid, [16](#)
  - snf\_handle\_t, [14](#)
  - snf\_init, [17](#)
  - snf\_link\_state, [14](#)
  - snf\_open, [17](#)
  - snf\_open\_defaults, [18](#)
  - snf\_ring\_close, [19](#)
  - snf\_ring\_getstats, [19](#)
  - snf\_ring\_open, [19](#)
  - snf\_ring\_open\_id, [20](#)
  - snf\_ring\_portinfo, [20](#)
  - snf\_ring\_rcv, [21](#)
  - snf\_ring\_rcv\_many, [21](#)
  - snf\_ring\_return\_many, [22](#)
  - snf\_ring\_t, [14](#)
  - snf\_set\_app\_id, [23](#)
  - snf\_start, [23](#)
  - snf\_stop, [23](#)
  - snf\_timesource\_state, [14](#)
- SNF\_F\_PSHARED
  - Open flags for process-sharing, port aggregation and packet duplication, [27](#)
- SNF\_RSS\_IPV4
  - Receive-Side Scaling (RSS), [25](#)
- SNF\_VERSION\_API
  - SNF API Reference, [13](#)
- snf, [37](#)
- snf\_close
  - SNF API Reference, [14](#)
- snf\_freeifaddrs
  - SNF API Reference, [15](#)
- snf\_get\_injection\_speed
  - Packet injection, [29](#)
- snf\_get\_link\_speed
  - SNF API Reference, [15](#)
- snf\_get\_link\_state
  - SNF API Reference, [15](#)
- snf\_get\_timesource\_state
  - SNF API Reference, [15](#)
- snf\_getifaddrs



- SNF API Reference, 16
- snf\_getportmask\_linkup
  - SNF API Reference, 16
- snf\_getportmask\_valid
  - SNF API Reference, 16
- snf\_handle\_t
  - SNF API Reference, 14
- snf\_ifa\_link\_speed
  - snf\_ifaddrs, 39
- snf\_ifa\_link\_state
  - snf\_ifaddrs, 39
- snf\_ifa\_macaddr
  - snf\_ifaddrs, 40
- snf\_ifa\_maxinject
  - snf\_ifaddrs, 40
- snf\_ifa\_maxrings
  - snf\_ifaddrs, 40
- snf\_ifa\_name
  - snf\_ifaddrs, 40
- snf\_ifa\_next
  - snf\_ifaddrs, 40
- snf\_ifa\_portnum
  - snf\_ifaddrs, 40
- snf\_ifaddrs, 39
  - pad, 39
  - snf\_ifa\_link\_speed, 39
  - snf\_ifa\_link\_state, 39
  - snf\_ifa\_macaddr, 40
  - snf\_ifa\_maxinject, 40
  - snf\_ifa\_maxrings, 40
  - snf\_ifa\_name, 40
  - snf\_ifa\_next, 40
  - snf\_ifa\_portnum, 40
- snf\_init
  - SNF API Reference, 17
- snf\_inject\_close
  - Packet injection, 29
- snf\_inject\_getstats
  - Packet injection, 29
- snf\_inject\_open
  - Packet injection, 30
- snf\_inject\_sched
  - Packet injection, 30
- snf\_inject\_sched\_v
  - Packet injection, 31
- snf\_inject\_send
  - Packet injection, 32
- snf\_inject\_send\_v
  - Packet injection, 33
- snf\_inject\_stats, 40
  - inj\_pkt\_send, 41
  - nic\_bytes\_send, 41
  - nic\_pkt\_send, 41
- snf\_inject\_t
  - Packet injection, 28
- snf\_link\_state
  - SNF API Reference, 14
- snf\_netdev\_reflect
  - Packet reflect to netdev (kernel stack), 35
- snf\_netdev\_reflect\_enable
  - Packet reflect to netdev (kernel stack), 36
- snf\_netdev\_reflect\_t
  - Packet reflect to netdev (kernel stack), 35
- snf\_open
  - SNF API Reference, 17
- snf\_open\_defaults
  - SNF API Reference, 18
- snf\_pkt\_fragment, 41
  - length, 41
  - ptr, 41
- snf\_pkt\_overflow
  - snf\_ring\_stats, 45
- snf\_recv\_req, 41
  - hw\_hash, 42
  - length, 42
  - length\_data, 42
  - pkt\_addr, 42
  - portnum, 42
  - timestamp, 42
- snf\_ring\_close
  - SNF API Reference, 19
- snf\_ring\_getstats
  - SNF API Reference, 19
- snf\_ring\_open
  - SNF API Reference, 19
- snf\_ring\_open\_id
  - SNF API Reference, 20
- snf\_ring\_portinfo, 42
  - data\_addr, 43
  - data\_size, 43
  - portcnt, 43
  - portmask, 43
  - q\_size, 43
  - ring, 43
  - SNF API Reference, 20
- snf\_ring\_qinfo, 43
  - q\_avail, 44
  - q\_borrowed, 44
  - q\_free, 44
- snf\_ring\_recv

- SNF API Reference, [21](#)
- snf\_ring\_recv\_many
  - SNF API Reference, [21](#)
- snf\_ring\_return\_many
  - SNF API Reference, [22](#)
- snf\_ring\_stats, [44](#)
  - nic\_bytes\_recv, [44](#)
  - nic\_pkt\_bad, [44](#)
  - nic\_pkt\_dropped, [45](#)
  - nic\_pkt\_overflow, [45](#)
  - nic\_pkt\_recv, [45](#)
  - ring\_pkt\_overflow, [45](#)
  - ring\_pkt\_recv, [45](#)
  - snf\_pkt\_overflow, [45](#)
- snf\_ring\_t
  - SNF API Reference, [14](#)
- snf\_rss\_mode\_flags
  - Receive-Side Scaling (RSS), [25](#)
- snf\_rss\_mode\_function, [45](#)
  - rss\_context, [46](#)
  - rss\_hash\_fn, [46](#)
- snf\_rss\_params, [47](#)
  - mode, [47](#)
  - params, [47](#)
  - rss\_flags, [47](#)
  - rss\_function, [47](#)
- snf\_rss\_params\_mode
  - Receive-Side Scaling (RSS), [26](#)
- snf\_set\_app\_id
  - SNF API Reference, [23](#)
- snf\_start
  - SNF API Reference, [23](#)
- snf\_stop
  - SNF API Reference, [23](#)
- snf\_timesource\_state
  - SNF API Reference, [14](#)
- timestamp
  - snf\_recv\_req, [42](#)