Introduction

The document is intended as a user's guide for the PacketMAX AF10G4ACE. It describes the hardware, value-add packet processing functions supported, operating instructions, and the user interface.

AF10G4ACE Hardware

As shown in Figure 1, the AF10G4ACE is a small form factor appliance, complete with an x86 CPU subsystem, that performs advanced packet conditioning applications. The AF10G4ACE design is illustrated in Figure 1. The main building blocks of the AF10G4ACE are the FPGA Motherboard, x86 ComExpress Modules, and a network Interface featuring a 4 x10G SFP+ frontend, and a 40G QSFP . The Precise Time Stamping source is an externally sourced 1 PPS. The unit features a Rugged Compact Design (1U x 8.25" x 14"). The AF10G4ACE faceplate is shown in Figure 2.

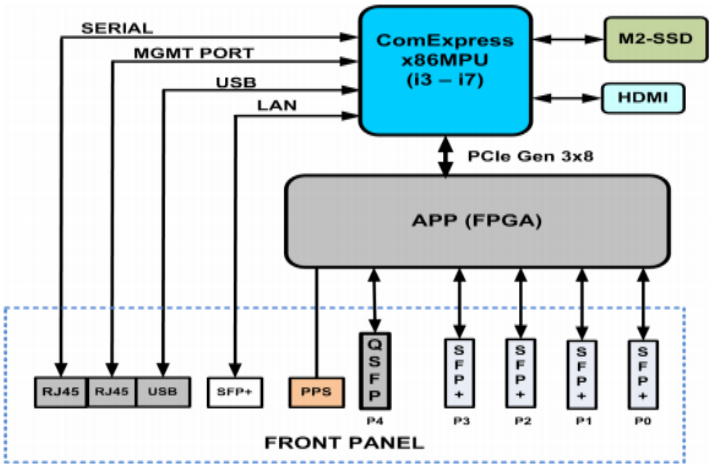

Figure 1. AF10G4ACE Hardware Design

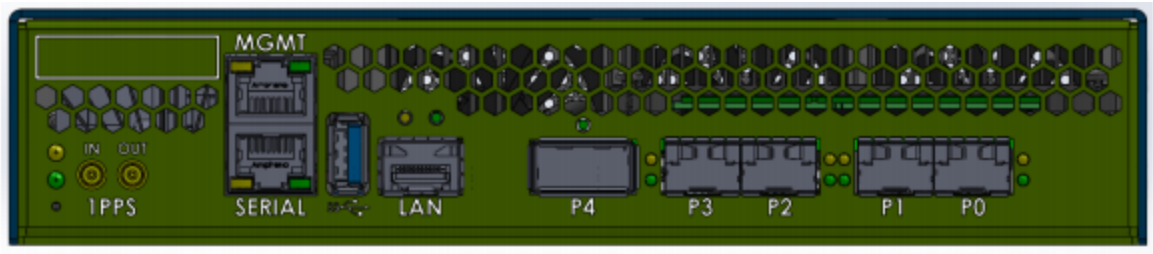

Figure 2. AF10G4ACE Front Panel

AF10G4ACE Service Node

The AF10G4ACE Service Node connects to device ports that send it streams of packets for packet conditioning on 10G or 40G Ethernet ports. The AF10G4ACE performs the required packet conditioning function on each ingress stream and performs a U-Turn on each link and returns the stream in the egress direction on the same device port. The AF10G4ACE can perform value-add functions that include deduplication, configurable Parsing, Time Stamping, Header Stripping and Packet Slicing. These value-add functions are described in more detail in the following sections.

Deduplication

Packet duplication is a potential problem in all large data center networks. While there are many causes of packet duplication, a common cause is networks that are tapped at multiple locations which tend to output duplicate copies of packets. Packet Brokers (PBs) or Network monitoring switches hence receive multiple copies of the same packet from these tapping points as the packets traverse the network. Removal of duplicate packets in the PB will therefore significantly reduce processing overhead in Monitoring Tools and Probes.

The AF10G4ACE is designed to remove all duplicate packets in the ingress stream within a roughly 850mS window, for an average packet size of 128 Bytes and send the de-duplicated packet stream back out the egress / Tx direction to the Packet Broker. The deduplication function is performed by the Advanced Packet Processor (APP), in conjunction with on-board Buffer Memory and Look-up Tables. Removal of duplicate packets reduces significant processing overhead in the Monitoring Applications running on a Network Probe connected to a Packet Broker.

Packet Parsing and Masking

In order for the AF10G4ACE to perform accurate de-duplication, the AF10G4ACE parser skips encapsulating headers to isolate the actual IP packet of interest and masks volatile IP header fields.

The sequence is as follows:

a L2 tags such as VLAN , MPLS and FabricPath are skipped by default. However, VLAN and MPLS may be optionally included if desired. The standard L2 header (destination MAC, source MAC, ethernet type) is always skipped. Optionally, support for ignoring VN-Tag can also be configured.

b The optional encapsulations that can be enabled include: IP-in-IP, GRE, GTP-U, VXLAN. (when enabled the inner IP frame is used for determining uniqueness).

At this point, the packet of interest has been isolated, starting with its L3 header.

The following header fields are masked (cleared to zero):

IPV4 TTL and header checksum

IPV6 hop limit

Two packets in the dedupe window are considered to be duplicates if after the above processing, they have the same hash signature.

Deduplication Performance:

The Deduplication process uses a hash-based table in order to store past packet signatures, and compares the current packet signature to past packet signatures to find a duplicate. There are two performance metrics:

1. Packets per second, which is determined by access rate into the DRAM table. This is limited to 30M packets per second. This translates into all four 10G ports or one 40G carrying full load traffic with an average packet size of 128 bytes.
2. Maximum deduplication window depth, which is determined by the DRAM table size, which is 16M buckets of depth 8 each. A probabilistic analysis shows a below 0.1% missed duplicate rate with about 30M packets in progress. With 4 ports of 10G, or 1 port of 40G, at 128B average packet size, this translates to a 850 ms dedupe window.

Time Stamping

The AF10G4ACE timestamps all packets at the ingress. The timestamp is internally a 64-bit integer, with 32-bits of seconds and 32-bits of sub-seconds. The resolution is that of the FPGA main clock, namely 5ns. The timestamp is disciplined using NTP by default, and using a front-panel 1pps input as an option. The 1pps can be derived from the usual GPS, PTP or CDMA sources.

This timestamp is carried together with the packet through the processing pipeline, and can optionally be added to egress packets in the Gigamon Header Time Stamp format.

Packet Brokers support proprietary time stamp formats, since there is no standard time stamping format. Because of its programmability, future releases of the AF10G4ACE may be programmed to emulate other specific time stamping formats.

Packet Slicing

If the target application in a monitoring tool does not require analysis of all packet data in its entirety, packet slicing can be used to limit the amount of data that is U-turned on each port to the PB. Some example applications where this functionality is useful are traffic engineering, billing or protocol analysis.

The AF10G4ACE packet slicing mode can be configured for two options:

a. Offset a programmable number of bytes from beginning of the Ethernet frame.
b. Offset a programmable number of bytes, anchored to various packet headers, namely the L2, L3 or L4 header.

1 Main Menu

The AF10G4ACE supports the following applications.

```
0: menu exit
1: deduplication control
2: slicing control
3: parser control
4: egress insertion control
5: egress stripping control
6: stats
7: configuration
8: restart application
9: reboot/shutdown appliance
10: interface mode
11: timing mode
```

1. Enter 1-11 (rt) for the desired application.

2 Interface Mode

The AF10G4ACE supports two interface modes, 4x10G or 40G. The default is 4x10G.

1. On the Main Menu, enter 10 (rt).

The Interface Mode Menu will be displayed.

```
WARNING: These options are only configured on application start.
Changes must be saved in the configuration and the application restarted to take
effect.

0: menu exit (default)
1: 4x10G
2: 40G

enter selection:
```

2. Enter 1 (rt) or 2 (rt).

3. Enter 0 (rt).

4. On the Main Menu, enter 7 (rt).

5. On the Configuration Menu, enter 2 (rt).

6. Enter 0 (rt).

7. On the Main Menu, enter 8 (rt).

3 Display the Configuration

The configuration of the AF10G4ACE is displayed under the Configuration Menu.

1. On the Main Menu, enter 7 (rt).

2. On the Configuration Menu, enter 1 (rt).

The configuration will be displayed.

```
dedup on 0
timestamp off
slice off
ipinip off
gre off
gtpu off
vntag off
vxlan off
egressvlan off
stripvlan off
stripmpls off
stripvntag off
stripfp off
stripvxlan off
stripgtpu off
stripipinip off
stripl3gre off
stripl2gre off
ifmode 40G
timingmode internal
```

4 Deduplication Control

1. On the Main Menu, enter 1 (rt).

The Deduplication Control Menu will be displayed.

```
0: menu exit (default)
1: disable deduplication
2: enable  deduplication

enter selection:
```

2. Enter 1 (rt) or 2 (rt).

3. If 2 was entered, the Validation Window Value Menu will appear.

```
1:    66.7 milliseconds (minimum, default)
2:   133.3 milliseconds
3:   200.0 milliseconds
4:   266.7 milliseconds
5:   333.3 milliseconds
6:   400.0 milliseconds
7:   466.7 milliseconds
8:   533.3 milliseconds
9:   600.0 milliseconds
10:  666.7 milliseconds
11:  733.3 milliseconds
12:  800.0 milliseconds
13:  866.7 milliseconds
14:  933.3 milliseconds
15: 1000.0 milliseconds (maximum)

enter window selection:
```

4. Enter 1-15 (rt) for the desired validation window value.

5. The VLAN tag support option will appear.

```
1: ignore VLAN tags in packet signature (default)
2: include VLAN tags in packet signature
```

6. Enter 2 (rt) to include VLAN tags.

7. The MPLS tag support option will appear.

```
1: ignore MPLS tags in packet signature (default)
2: include MPLS tags in packet signature
```

8. Enter 2 (rt) to include MPLS tags.

9. Enter 0 (rt).

10. On the Main Menu, enter 7 (rt).

11. On the Configuration Menu, enter 2 (rt).

12. On the Configuration Menu, enter 0 (rt).

13. On the Main Menu, enter 8 (rt).

5 Slicing Control

Packet slicing may be controlled starting with the l2, l3 or l4 header.

1. On the Main Menu, enter 2 (rt).

The Slicing Control Menu will be displayed.

```
0: menu exit (default)
1: disable slicing
2: enable  slicing

enter selection:
```

2. Enter 1 (rt) or 2 (rt).

3. If 2 was entered:

the enter l2offset option will appear. Enter the number of bytes to slice from the l2 header. Enter -1 to disable. Disable if l3 or l4 slicing is desired.

```
enter l2offset value for non-IP packets (i.e. number of bytes after L2 header)
or:
-1: disable slicing for non-IP packets (default)

enter l2offset:
```

the enter l3offset option will appear. Enter the number of bytes to slice from the l3 header. Enter -1 to disable. If l2 header slicing was previously defined, enter -2. Disable if l4 slicing is desired.

```
enter l3offset value for IP packets other than TCP/UDP/SCTP (i.e. number of bytes
after
L3 header)
or:
-1: disable slicing for IP packets other than TCP/UDP/SCTP (default)
-2: use l2offset (i.e. number of bytes after L2 header)

enter l3offset:
```

the enter l4offset option will appear. Enter the number of bytes to slice from the l4 header. Enter -1 to disable. If l2 or l3 header slicing was previously defined, enter -2 or -3.

```
enter l4offset value for TCP/UDP/SCTP packets (i.e. number of bytes after L4
header)
or:
-1: disable slicing for TCP/UDP/SCTP packets (default)
-2: use l2offset (i.e. number of bytes after L2 header)
-3: use l3offset (i.e. number of bytes after L3 header)

enter l4offset:
```

4. Enter 0 (rt).

5. On the Main Menu, enter 7 (rt).

6. On the Configuration Menu, enter 2 (rt).

6 Parser Control

1. On the Main Menu, enter 3 (rt).

The Parser Control Menu will be displayed.

```
WARNING: These options are only configured on application start.
Changes must be saved in the configuration and the application restarted to take
effect.

0: menu exit (default)
1: disable IP-in-IP encapsulation awareness
2: enable  IP-in-IP encapsulation awareness
3: disable GRE encapsulation awareness
4: enable  GRE encapsulation awareness
5: disable GTP-U encapsulation awareness
6: enable  GTP-U encapsulation awareness
7: disable VNTAG awareness
8: enable  VNTAG awareness
9: disable VXLAN encapsulation awareness
10: enable  VXLAN encapsulation awareness

enter selection:
```

2. Enter 1-10 (rt) for the desired option(s). More than 1 option may be enabled/disabled.

3. Enter 0 (rt).

4. On the Main Menu, enter 7 (rt).

5. On the Configuration Menu, enter 2 (rt).

6. Enter 0 (rt).

7. On the Main Menu, enter 8 (rt).

7 Egress Insertion Control

1. On the Main Menu, enter 4 (rt).

The Egress Insertion Menu will be displayed.

```
0: menu exit (default)
1: disable timestamp insertion
2: enable  timestamp insertion
```

```
3: disable egress VLAN insertion
4: enable  egress VLAN insertion

enter selection:
```

2. Enter 1-2 (rt) to enable/disable timestamp insertion.

3. Enter 3-4 (rt) to enable/disable egress VLAN insertion.

4. If 4 was entered the unit will prompt for the VLANs for each port.

```
enter port 0 VLAN tag (default 100):
enter port 1 VLAN tag (default 101):
enter port 2 VLAN tag (default 102):
enter port 3 VLAN tag (default 103):
```

5. Enter 0 (rt).

6. On the Main Menu, enter 7 (rt).

7. On the Configuration Menu, enter 2 (rt).

8 Egress Stripping Control

1. On the Main Menu, enter 5 (rt).

The Egress Stripping Menu will be displayed.

```
WARNING: These options are only configured on application start.
Changes must be saved in the configuration and the application restarted to take
 effect.


0: menu exit (default)
1: disable VLAN stripping
2: enable  VLAN stripping
3: disable MPLS stripping
4: enable  MPLS stripping
5: disable VNTAG stripping
6: enable  VNTAG stripping
7: disable FabricPath stripping
8: enable  FabricPath stripping
9: disable VXLAN stripping
10: enable  VXLAN stripping
11: disable GTP-U stripping
12: enable  GTP-U stripping
13: disable IP-in-IP stripping
14: enable  IP-in-IP stripping
15: disable L3GRE stripping
16: enable  L3GRE stripping
```

```
17: disable L2GRE stripping
18: enable  L2GRE stripping

enter selection:
```

2. Enter 1-18 (rt) for the desired option(s). More than 1 option may be enabled/disabled.

3. Enter 0 (rt).

4. On the Main Menu, enter 7 (rt).

5. On the Configuration Menu, enter 2 (rt).

6. Enter 0 (rt).

7. On the Main Menu, enter 8 (rt).

9 Stats

1. On the Main Menu, enter 6 (rt).

The stats for the ingress and egress ports will be displayed. The stats refresh 1/sec.

```
Elap    6 (0:00:06)   16:04:53   2020-10-05
ingress port  packets   bytes   malfs    rsrcs     fifo%     1sec b/w
port      0      0        0       0        0       0.000     0.000     DOWN
port      1      0        0       0        0       0.000     0.000     DOWN
port      2      0        0       0        0       0.000     0.000     DOWN
port      3      0        0       0        0       0.000     0.000     DOWN
port      4      0        0       0        0       0.000     0.000     DOWN
ingress total    0        0       0        0                 0.000
egress  port  packets   bytes
egress    0      0        0
egress    1      0        0
egress    2      0        0
egress    3      0        0
egress    4      0        0
egress  total    0
```

2. Press the Enter key to halt the stats and return to the Main Menu.

The stats are divided into two categories, ingress and egress.

Ingress Stats

```
ingress port  packets   bytes   malfs    rsrcs     fifo%     1sec b/w
port      0      0        0       0        0       0.000     0.000     DOWN
port      1      0        0       0        0       0.000     0.000     DOWN
port      2      0        0       0        0       0.000     0.000     DOWN
port      3      0        0       0        0       0.000     0.000     DOWN
```

```
port      4      0        0        0        0       0.000     0.000     DOWN
ingress total    0        0        0        0                 0.000
```

port 0	10G port when the unit is in the 4x10G interface mode.
port 1	10G port when the unit is in the 4x10G interface mode.
port 2	10G port when the unit is in the 4x10G interface mode.
port 3	10G port when the unit is in the 4x10G interface mode.
port 4	40G port when the unit is in the 40G interface mode.
ingress total	stats for all ingress ports

packets	number of packets received
bytes	number of bytes received
malfs	number of malformed packets received
rsrcs	number of packets dropped due to the queue is full
fifo%	queue percentage being used
1sec b/w	port bandwidth

Egress Stats

```
egress  port  packets   bytes
egress    0      0        0
egress    1      0        0
egress    2      0        0
egress    3      0        0
egress    4      0        0
egress  total    0
```

port 0	10G port when the unit is in the 4x10G interface mode.
port 1	10G port when the unit is in the 4x10G interface mode.
port 2	10G port when the unit is in the 4x10G interface mode.
port 3	10G port when the unit is in the 4x10G interface mode.
port 4	40G port when the unit is in the 40G interface mode.
egress total	stats for all egress ports

packets	number of packets transmitted
bytes	number of bytes transmitted

10 Configuration

1. On the Main Menu, enter 7 (rt).

The Configuration Menu will be displayed.

```
0: menu exit (default)
1: display saved configuration
```

```
2: save configuration
3: clear configuration
7: print version info
9: network configuration

enter selection:
```

2. Enter 0 (rt) to return to the Main Menu.

3. Enter 1 (rt) to display the saved configuration.

4. Enter 2 (rt) to save the configuration.

5. Enter 3 (rt) to clear the configuration to the default. This will not reset the network configuration.

6. Enter 7 (rt) to print the AF10G4ACE version information.

```
SDK:      SDK_1_2_20210928
firmware: 49410c03
```

7. Enter 9 (rt) to setup the network configuration of the AF10G4ACE. This process is discussed in the AF10G4ACE Initial Setup Guide.

11 Timing Mode

1. On the Main Menu, enter 11 (rt).

The Timing Mode Menu will be displayed.

```
WARNING: These options are only configured on application start.
Changes must be saved in the configuration and the application restarted to take
effect.

0: menu exit (default)
1: internal appliance clock (NTP disciplined)
2: PPS rising edge
3: PPS falling edge

enter selection:
```

2. Enter the desired timing mode option, 1,2,3 (rt).

3. Enter 0 (rt).

4. On the Main Menu, enter 7 (rt).

5. On the Configuration Menu, enter 2 (rt).

6. Enter 0 (rt).

7. On the Main Menu, enter 8 (rt).

12 Restart Application

If applications require a restart to take affect the following warning message will be displayed on the menu for
that application.

```
WARNING: These options are only configured on application start.
Changes must be saved in the configuration and the application restarted to take
effect.
```

1. On the Main Menu, enter 8 (rt).

13 Reboot/Shutdown Appliance

1. On the Main Menu, enter 9 (rt).

The Reboot/Shutdown Menu will be displayed.

```
0: menu exit (default)
1: reboot
2: graceful shutdown

enter selection:
```

2. Enter 1 (rt) to reboot the AF10G4ACE.

3. Enter 2 (rt) to gracefully shutdown the AF10G4ACE. A graceful shutdown should always be executed before the AF10G4ACE is turned off or AC power removed.