NetApp Katana Pensando User Guide

Version: 0.12.0
Date: July 9, 2019

1 Pensando Legal Note

All information in this document is provided on a non-disclosure basis. Anyone reading this document implicitly agrees to be bound by Pensando Systems' non-disclosure terms.

Document Glossary

The following table describes the product "code names" used in this document and release. Final product names are TBD.

Name	Reference
NAPLES	Pensando Systems Adapter Card. ""Network Adapter with Programmable Logic and Enhanced Security/Storage/Services"
Capri	Pensando System custom ASIC on the NAPLES Card
SONIC	Pensando Kernel driver that supports compress/decompression services in hardware
IONIC	Pensando Kernel driver for NAPLES network adapter
penctl	CLI utility for managing NAPLES from the local host

2 NAPLES Product Overview

The NAPLES Product is designed to connect rapidly evolving compute and storage servers to a new generation of scalable, programmable data center networks. The NAPLES product includes a highly configurable network pipeline, a customizable host interface, and flexible hardware offload engines for storage, security, and network functions that can be processed in a pipelined manner at line rate.

The name “NAPLES” stands for “Network Adapter with Programmable Logic and Enhanced Security/Storage/Services”

The NAPLES product is based on a customized ASIC (“Capri”) that implements a full match + action pipeline. All traffic is fully programmable through P4 network programming specification.

The Capri ASIC implements a match + action pipeline using parsers, deparsers, table engines and match processing units (MPU). The forwarding logic is expressed in P4 language as a P4 program.

The NAPLES software subsystem runs on the ARM cores present in the Capri ASIC and includes Compression/Decompression offload as provided through the SONIC driver.

This document describes the NAPLES product. Tasks for managing the control and data plane of Naples will be explained in this document. Internals of Naples card and offload architecture will be shown. Administrators should use this guide to understand how to use and manage Naples.

3 Naples Components

Naples is a PCIe card providing for offload of network, storage, and security services. Naples is a full-height half-width PCIe card containing a processor (Capri), network ports, management port and flash memory.

3.1 Naples Hardware Diagram

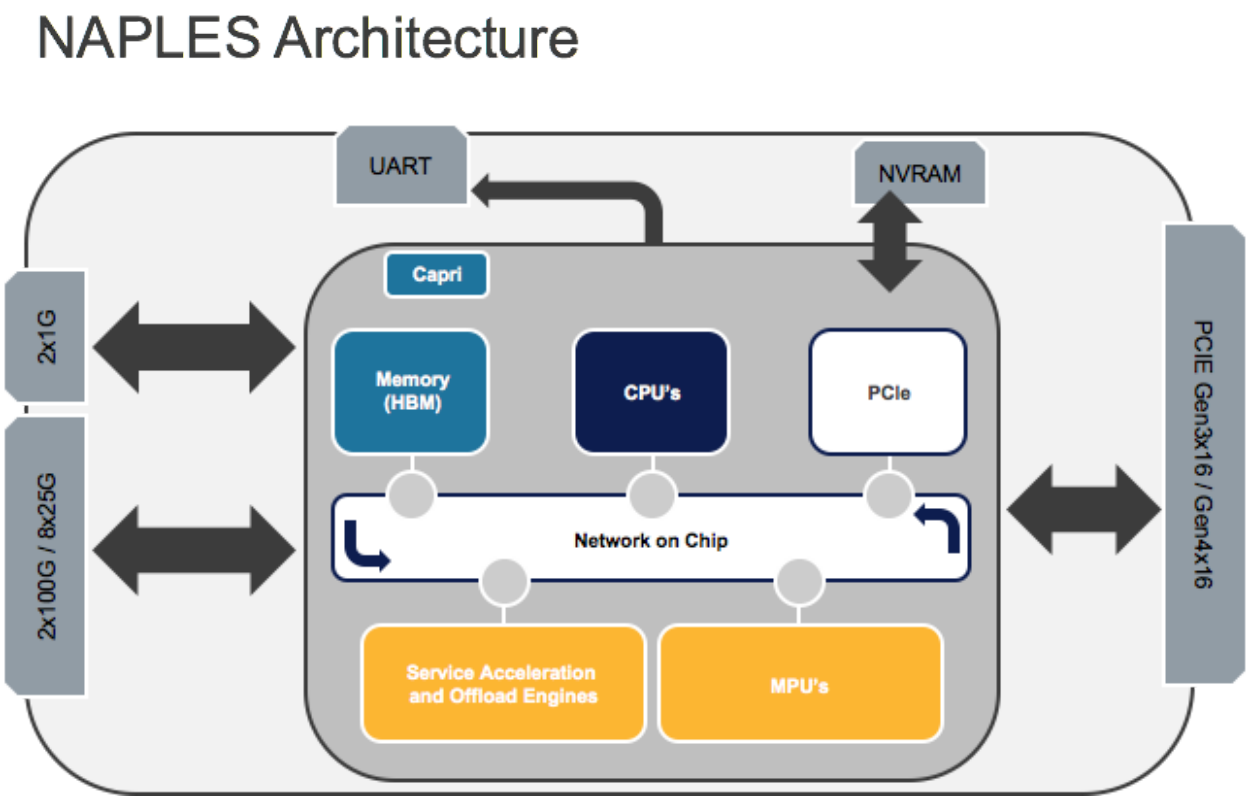

The “2x100G / 8x25G” uplinks label refers to the various link speeds that are supported, including:

- 2x 100G
- 4x 50G
- 4x 40G
- 8x 25G

3.2 Naples / Capri Overview

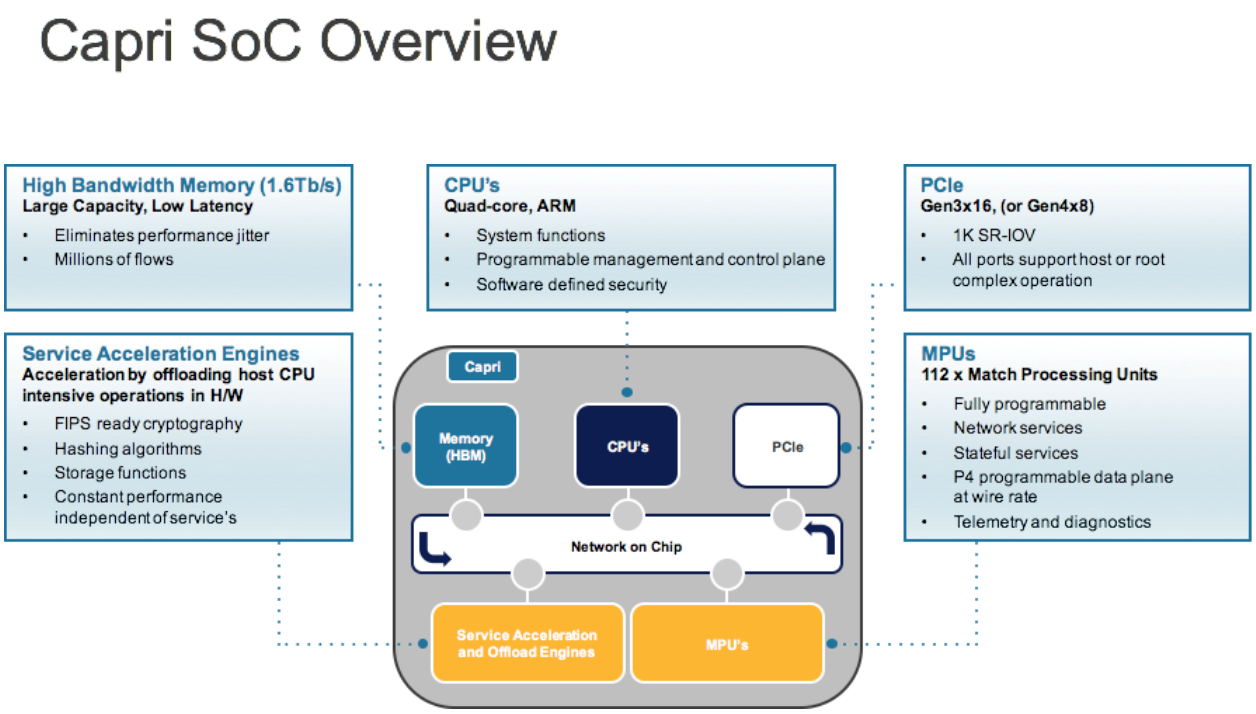

3.3 NAPLES Software and Services

NAPLES achieves service offload by implementing full network stack on its custom ASIC chip Capri. As presented in the above diagram, the platform contains memory to support data path operations, ARM CPUs to support the control path, and interfaces towards network and PCIe. To process packets according to programmed network policy, match processing units (MPUs) are used. Service acceleration is used to accelerate encryption, compression, hashing and storage functions. This highly programmable data path, with full network stack on the NAPLES product provides for hardware accelerated, pipelined services offload at wire speed.

4 Installing Naples into the Server

4.1 Adapter Card Installation Instructions

Installation of the Naples adapter card requires following standard safety procedures for working with systems sensitive to static electricity discharge. The following safety procedures involve:

1. Removing any metallic objects from hands and wrists.
2. Making sure to use only insulated tools as shown on the picture below.
3. Verifying that the system is powered off and is unplugged.
4. It is strongly recommended to use an ESD strap shown on a picture below or other antistatic devices.

To successfully install the card, double check the following:

1. Verify that the system meets the hardware and software requirements
2. After shutting down the system, turn off power and unplug the cord.
3. Remove the card from its package. Please note that the card must be placed on an antistatic surface.
4. Check the card for visible signs of damage. Do not attempt to install the card if damaged.

To Install the card, follow the instructions below:

1. Before installing the card, make sure that the system is off and the power cord is not connected to the server. Please follow proper electrical grounding procedures.
2. Open the system case.
3. Locate an available PCI Express slot for the adapter card. Please avoid damaging the LEDs with a screwdriver. Do not force the bracket onto the adapter card as to not damage the EMI fingers on the cages. Please note that the following figures are for

illustration purposes only. A lesser width adapter can be seated into a greater width slot (x8 in a x16), but a greater width adapter cannot be seated into a lesser width slot (x16 in a x8). Align the adapter connector edge with the PCI Express connector slot.
4. Applying even pressure at both corners of the card, insert the adapter card into the PCI Express slot until firmly seated.
5. When the adapter is properly seated, the port connectors are aligned with the slot opening, and the adapter faceplate is visible against the system chassis.

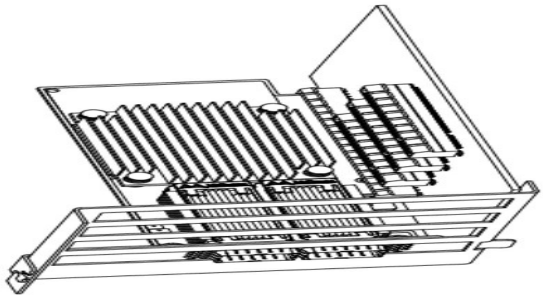

6. Secure the adapter with the adapter clips or screws.
7. Close the system case.

4.2 Adapter Card Un-installation Instructions

Use Safety Precautions when working on removing the network card. Follow the safety instructions described above for installing the network card. The adapter is installed in a system that operates at high voltages.

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. To remove the card, disengage the retention mechanisms on the bracket (clips or screws).
4. Holding the adapter card from its center, gently pull the adapter card from the PCI Express slot.
5. When the port connectors reach the top of the chassis window, gently pull the adapter card in parallel to the motherboard like shown on Figure 1.

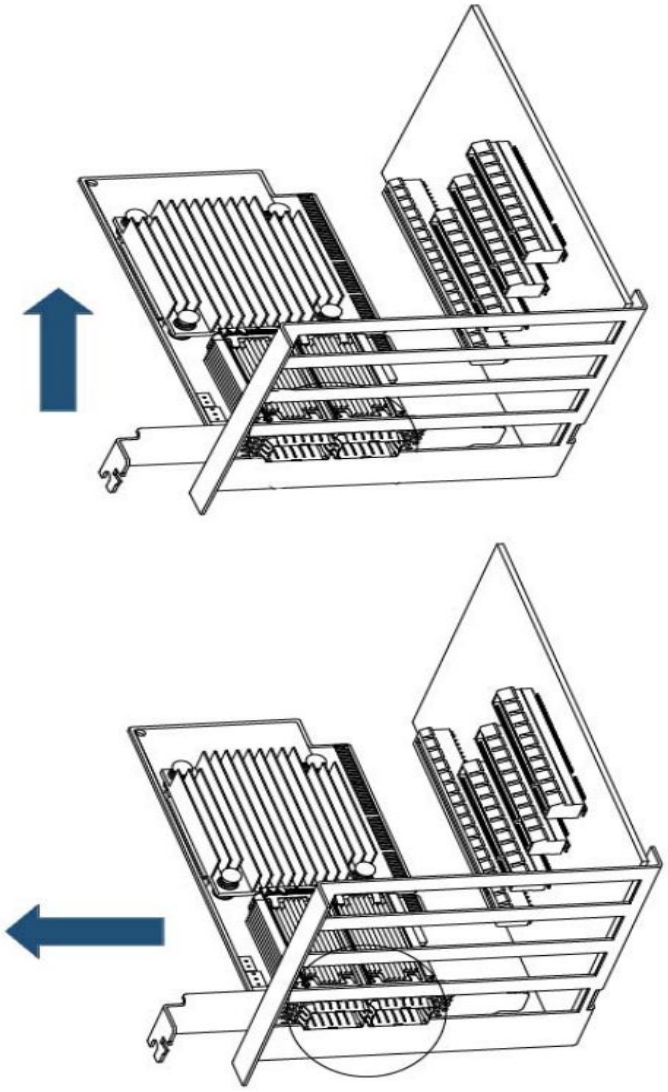

.

4.3 PCIe Supported

The Naples Edge Services Module requires hardware systems that support PCI Express Gen 3 (1.0 and 2.0 compatible) x16 slots. The device can be either an initiator/master, initiating the PCI Express bus operations, or a target/slave, responding to PCI bus operations.
The following lists PCIe interface features:

- PCIe Gen 3.0 compliant, 1.1 and 2.0 compatible
- 2.5, 5.0, or 8.0 GT/s link rate x8 and x16
- Auto-negotiates to x16, x8, x4, x2, or x1

- Support for MSI/MSI-X mechanisms
- Advanced Error Reporting (AER) capability
- PCIe Atomic operations
- SRIOV up to 1000 virtual functions

4.4 Number of Naples Supported

The number of Naples Edge Services Modules depend on the number of PCIe slots on the motherboard and the server's available power budget.

4.5 Power and Cooling Requirements

Naples Edge Services Module requires that the server's cooling system support cooling of a maximum of 50W during regular operations. Specifications can be found in Section 7 of this document.

5 Naples **penctl** Configuration Utility

penctl is the management CLI used to manage NAPLES from the host, when in the default "host-managed" mode.

Common Use Cases for **penctl**

- Naples F/W Management**:** install f/w bundles, view current f/w versions, set f/w version for subsequent reboots.
- Change from "host-managed" (default) to "network-secure" mode**.** In "network-secure" management mode, Naples policies cannot be configured or modified from the host.
  NB: Once Naples has been successfully changed by establishing a connection to **Venice** (or another external controller), only the external controller can change back to "host-managed" mode.
- Naples Logs, Events and System Information**:** allows examining the following:
  System logs (e.g. coming from Linux kernel, boot logs)
  Process logs (e.g. debug output from various processes that may be running in Naples)
  Alerts, warnings produced by the system
  System information: flash disk usage, cpu usage, memory usage etc.
  Hardware health: parameters e.g. temperature, datapath health, etc.

When NAPLES presents to the host, there are 3 network interfaces, as shown below:

```
ionic0: flags=8802<BROADCAST,SIMPLEX,MULTICAST> metric 0 mtu 1500
options=e507bb<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING,JUMBO_MTU,VLAN_HWCSUM,TSO4
,TSO6,LRO,VLAN_HWFILTER,VLAN_HWTSO,RXCSUM_IPV6,TXCSUM_IPV6>
       ether 00:de:ad:be:ef:02
       hwaddr 00:de:ad:be:ef:02
       nd6 options=29<PERFORMNUD,IFDISABLED,AUTO_LINKLOCAL>
       media: Ethernet autoselect
       status: no carrier
ionic1: flags=8802<BROADCAST,SIMPLEX,MULTICAST> metric 0 mtu 1500
options=e507bb<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING,JUMBO_MTU,VLAN_HWCSUM,TSO4
,TSO6,LRO,VLAN_HWFILTER,VLAN_HWTSO,RXCSUM_IPV6,TXCSUM_IPV6>
       ether 00:de:ad:be:ef:03
       hwaddr 00:de:ad:be:ef:03
       nd6 options=29<PERFORMNUD,IFDISABLED,AUTO_LINKLOCAL>
       media: Ethernet autoselect
       status: no carrier
ionic2: flags=8802<BROADCAST,SIMPLEX,MULTICAST> metric 0 mtu 1500
options=e507bb<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING,JUMBO_MTU,VLAN_HWCSUM,TSO4
,TSO6,LRO,VLAN_HWFILTER,VLAN_HWTSO,RXCSUM_IPV6,TXCSUM_IPV6>
       ether 00:de:ad:be:ef:04
```

```
        hwaddr 00:de:ad:be:ef:04
        nd6 options=29<PERFORMNUD,IFDISABLED,AUTO_LINKLOCAL>
        media: Ethernet autoselect (1000Base-KX <full-duplex>)
        status: active
```

In the above output, `ionic0` and `ionic1` are the primary data network interfaces, and `ionic2` is reserved for NAPLES management.

For Linux-based distributions, the interfaces will present through `ip addr` as follows:

```
7: enp181s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UNKNOWN
group default qlen 1000
    link/ether 00:de:ad:be:ef:02 brd ff:ff:ff:ff:ff:ff
8: enp182s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UNKNOWN
group default qlen 1000
    link/ether 00:de:ad:be:ef:03 brd ff:ff:ff:ff:ff:ff
9: enp183s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
UNKNOWN group default qlen 1000
    link/ether 00:de:ad:be:ef:04 brd ff:ff:ff:ff:ff:ff
```

In the above example, `enp183s0` is the management interface.

The management interface is always presented as the 3rd interface.

The management interface requires a non-routable IP address assigned from the range **`169.254.0.[2-254]/24`**

penctl Installation and Setup FreeBSD and Linux

Requirements prior to installing **penctl**

- Make sure you have downloaded the appropriate distribution **tar** file available from Pensando Systems that include Naples F/W, **penctl** binary and driver sources for a particular OS type.
- **penctl** requires building and loading of the Naples driver for a given OS type.
- **penctl** communicates through the Naples management interface and requires that the management interface be configured with an IP address in the range 169.254.0.[2-254]/24.

Install penctl in an appropriate directory. Ex:

```
# cp penctl.linux /usr/local/bin/penctl    OR
# cp penctl.freebsd /usr/local/bin/penctl
```

Configure the NAPLES mgmt interface.
Linux Ex:

```
# PCIDEV=`lspci -d 1dd8: | grep Ethernet | sed -n 3p | awk '{print $1}'`
# MNIC=`lshw -c network -businfo | grep $PCIDEV  | awk '{print $2}'`
# ifconfig $MNIC 169.254.0.2/24
OR
# ip addr add 169.254.0.2/24 dev $MNIC
```

FreeBSD Ex:

```
# ifconfig ionic2 169.254.0.2/24
```

- **penctl** uses an environment variable called **`NAPLES_URL`** to communicate with the NAPLES card.

  Example:

  ```
  # export NAPLES_URL=http://169.254.0.1
  OR
  # setenv NAPLES_URL http://169.254.0.1
  ```

Using the **penctl** Command

The **penctl** Command has a broad variety of features, accompanied by extensive online help.

Please consult the online help and `--help` options for detailed usage and explanation.

Firmware Install Process

Naples card uses two locations to store firmware images: “**mainfwa**” and “**mainfwb**”. When Naples starts up, it will boot from one to the two firmware locations based on the “startup-firmware” setting.

Upon uploading a new firmware image via `penctl system firmware-install`, the “startup-firmware” pointer will be set to the new firmware image to run from next reboot.

A reboot is always required, when updating firmware.

The “startup-firmware” pointer can always be set manually via `update startup-firmware`

The current firmware versions and startup pointer can always be queried via:

```
penctl show firmware-version
```

Update Naples-profile to change default network interfaces status

By default Naples, uses a “default profile” where network interfaces are “UP” after power on. The behavior can be changed by applying a new naples “profile” so that network interfaces will be kept “DOWN” before driver loads.
Note: The profile/behavior change is persistent across host/naples reboot.

Below are the example steps to create and apply new napels “profiles”. Please see “penctl --help” to check all available options:

1. Create new naples-profile
   ```
   # penctl create naples-profile -n <profile-name> -p
   disable
   ```
2. Apply the profile to naples
   ```
   # penctl update naples -f <profile-name>
   ```
3. A host reboot is required to let newly applied naples-profile take effect.

6 NAPLES Troubleshooting

6.1 NAPLES Device is not detected by the host.

Address possible reasons:

- Bad PCIe connection (reseat the card in the PCIe slot)
- PCIe slot is disabled from the BIOS (check BIOS settings)
- Incorrect firmware image

Sample device presentation:

```
# lspci -d 1dd8:
5e:00.0 PCI bridge: Device 1dd8:1000
5f:00.0 PCI bridge: Device 1dd8:1001
5f:01.0 PCI bridge: Device 1dd8:1001
5f:02.0 PCI bridge: Device 1dd8:1001
5f:03.0 PCI bridge: Device 1dd8:1001
60:00.0 Ethernet controller: Device 1dd8:1002
61:00.0 Ethernet controller: Device 1dd8:1002
62:00.0 Ethernet controller: Device 1dd8:1004
63:00.0 Processing accelerators: Device 1dd8:1007
```

In the above example:

`60:00.0 Ethernet controller: Device 1dd8:1002`	NAPLES Ethernet Port 1
`61:00.0 Ethernet controller: Device 1dd8:1002`	NAPLES Ethernet Port 2
`62:00.0 Ethernet controller: Device 1dd8:1004`	NAPLES Management Port
`63:00.0 Processing accelerators: Device 1dd8:1007`	NAPLES Storage Offload

Also check “dmesg” for possible errors.

7 Specifications

7.1 <NAPLES-100 (OPN TBD)> Specifications

Table 5. <NAPLES-100 (OPN TBD)> Specification Table

Form factor and dimensions	**Size**	Full Height Half Length 4.53 in. x 5.60 in. (115.00mm x 142.24 mm)
	Connector	Dual QSFP28 (copper and optical)
Protocol Support	**Ethernet**	100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50G Ethernet Consortium, 40GBASE-CR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 25GBASE-CR/CR-S, 25GBASE-SR, 25GBASE-LR, 25G Ethernet Consortium, 10GBASE-CR, 10GBASE-SR, 10GBASE-LR, 10GBASE-ER
	Data Rate	10/25/40/50/100 Gb/s Ethernet
	PCI Express	Gen 3, SERDES @ 8.0GT/s, 16 lanes; Gen 4, SERDES @ 16.0GT/s, 8 lanes; Gen 2.0 and 1.1 compatible
Power and Environmental	**Voltage**	12V, 3.3V
	Power	**Cable Type**
	Typical Power	**Passive Cables**:
	Maximum Power	**Passive Cables**:
		1.5W Active Cables/Transceivers:
		2.5W Active Cables/Transceivers:
	Maximum power available through QSFP28 port: 3.5W	
	Temperature	**Operational:** 0°C to 55°C
		Non-operational: -40°C to 70°C
	Humidity	90% relative humidity
	Air Flow	See Airflow Specifications on page xxx
Regulatory	**Safety**	CB / cTUVus / CE (to be confirmed)

	EMC	CE / FCC
	ROHS	RoHS-R6

7.2 Airflow Specifications

Table 7. Airflow Specification Table [1]

Air Flow (LFM) - TBD					
Cable Type	Passive	Active 1.5W	Active 2.5W	Active 3.5W	Active 5W
NAPLES-100	TBD	TBD	TBD	TBD	TBD

1) Air Flow Direction - Heat Sink to Port