ET2500 Edge Network Appliance

User Manual

Product Model	ET2500
Product Version	V2.0
Document Version	V2.0
Intended Audience	User/Maintainer/Developer

Copyright ©2026 Asterfusion. All rights reserved.

Contents

Copyright ©2026 Asterfusion. All rights reserved.

Copyright ©2026 Asterfusion. All rights reserved.

Copyright ©2026 Asterfusion. All rights reserved.

About This Document:

This user manual introduces the Asterfusion ET2500 edge network appliance, providing information on the hardware parameters and specifications related to the network appliance, as well as details about the required software and firmware for the network appliance. It also offers step-by-step instructions on how to install the operating system on the network appliance to ensure the device functions properly.

The final section of the document includes detailed testing methods and procedures for conducting DPDK traffic tests and VPP performance tests based on the ET2500.

Intended Audience:

This user manual is intended for ET2500 edge network appliance users, installers, and developers.

Technical Support:

EMAIL: support@asteraix.com

Revision History:

Revision Time	Revision Content	Note
2024/10/09	Create	V1.0
2026/07/10	Update	V2.0

Copyright ©2026 Asterfusion. All rights reserved.

1 Preface

1.1 Reference

Serial Number	Reference Documents
1	https://www.dpdk.org/
2	https://s3-docs.fd.io/vpp/24.02/
3	https://github.com/asterfusion/Helium_DPU/

Table 1-1 References

1.2 Abbreviations

Abbreviations	EN	CN
PCIe	Peripheral Component Interconnect Express	高速串行计算机扩展总线标准
DPDK	Data Plane Development Kit	数据平面开发套件
VPP	Vector Packet Processor	

Table 1-2 Abbreviations

1.3 Command Description

In this document, all commands displayed in gray indicate that they can be run directly on the device.

The sample is as follows:

```
# uname -a
```

1.4 Naming Specifications for Helium Product Packages

The naming specifications of the ET2500 software package are as follows:

Product - Function – Version

The explanations are as follows, where parentheses indicate non-essential items:

1. Product: ET2500
2. Function: Describes the functionality provided or available by the package
3. Version : Release version of the software package. The version format is Vx.yRz-(date)
    a) x - major version number, used for incompatibility changes
    b) y - minor version number, used to identify new or modified features
    c) z - revised version, used to identify problem fixes

Copyright ©2026 Asterfusion. All rights reserved.

d) date - the release date corresponding to the software package

Copyright ©2026 Asterfusion. All rights reserved.

2 Introduction

2.1 Product Overview

The Asterusion ET2500 is an edge network appliance designed to address the intricate networking requirements of modern enterprises. In recent years, businesses have embraced technologies like cloud computing, big data, video conferencing, live streaming, and notably, the application of AI and IoT technologies has facilitated business transformation through machine empowerment. Traditional enterprise networks typically employ various dedicated devices to handle different functions, making enterprise networks complex, expensive, and difficult to operation & maintain. To address this issue, the ET2500 adopts a Computing-Networking fusion chip architecture and a decoupled open design, allowing enterprises to use a single intelligent device to perform all of the aforementioned functions.

2.2 Physical Drawing of ET2500

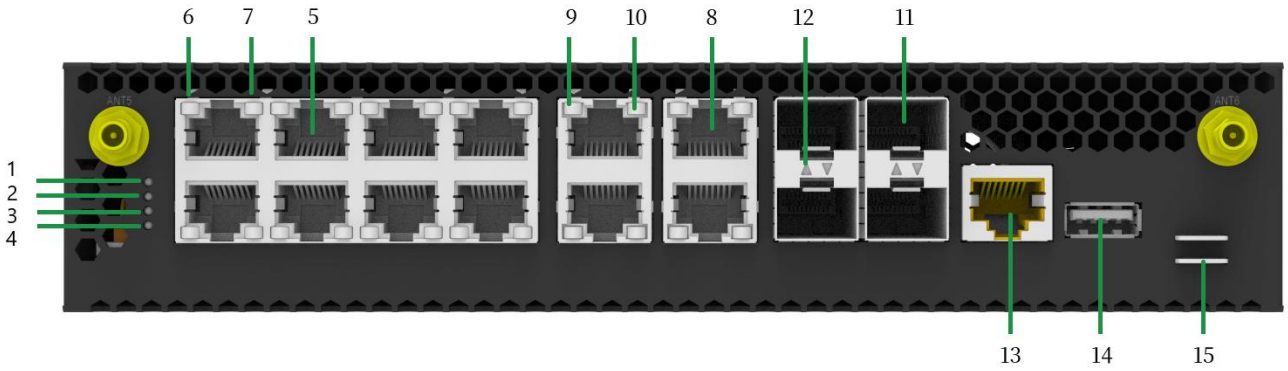

Figure 2-1 ET2500 Physical Products Front

Figure 2-2 ET2500 Physical Products Rear

The following table lists the meanings of the notes on the actual drawings of the ET2500.

Copyright ©2026 Asterfusion. All rights reserved.

Label	Explanation
Label 1	Power LED
Label 2	System LED
Label 3	Fan LED
Label 4	Location LED
Label 5	1G RJ45 Port
Label 6	1G RJ45 Port Left LED
Label 7	1G RJ45 Port Right LED
Label 8	2.5G RJ45 Port
Label 9	2.5G RJ45 Port Left LED
Label 10	2.5G RJ45 Port Right LED
Label 11	10G SFP+ Port
Label 12	10G SFP+ Port LED
Label 13	Serial Port
Label 14	USB Port
Label 15	SIM

Table 2-1 Label Meaning Table of Physical Products Front

Label	Explanation
Label 1	Fan 1
Label 2	Fan 2
Label 3	Power

Table 2-2 Label Meaning Table of Physical Products Rear

Front Panel Status Port LED		
LED	Status	Description
Power LED(Power)	Green, always on	Power ok
	Red, always on	Power fail
Sys LED(S)	Green, always on	System is on and ready
	Green, blink	Booting
	Red, always on	When detect HW failed
	OFF	No power
Fan LED(F)	Green, always on	All fans are working on normal
	Red, always on	Any fan tray fault
	OFF	No power
Location LED(L)	Blue, blink in 4Hz	Set by management to locate switch.

Copyright ©2026 Asterfusion. All rights reserved.

		Use command led_loaction to active.
	OFF	Function not active. Use command led_location_off to inactive.

Table 2-3 Front Panel Status Port LED

Front Panel Service Port LED		
LED	Status	Description
SFP+ LED	Green, always on	Link up, no data transmission
	Green, Blink	Link up, data transmission
	OFF	Link down
RJ45 LNK/ACT LED	Orange, always on	Link up, no data transmission
	Orange, Blink	Link up, data transmission
	OFF	Link down
RJ45 PoE LED	Green, always on	PoE port is supplying power for downlink port
	OFF	No PoE Power supply
RJ45 GIGA LED	Green, always on	1G/2.5G Link
	OFF	Link down or 10M/100M link

Table 2-4 Front Panel Service Port LED

Copyright ©2026 Asterfusion. All rights reserved.

3 Installation Process

3.1 Scope of Application

Scene	Description	Reference
Install operating system		8 Appendix A Operating System Installation
Set the port rate		5.3.2 Port Rate Settings for Panel Ports
DPDK runtime test		6 DPDK Runtime Test
VPP runtime test		7 VPP Runtime Test

Table 3-1 Scope of Application

3.2 Installation Steps

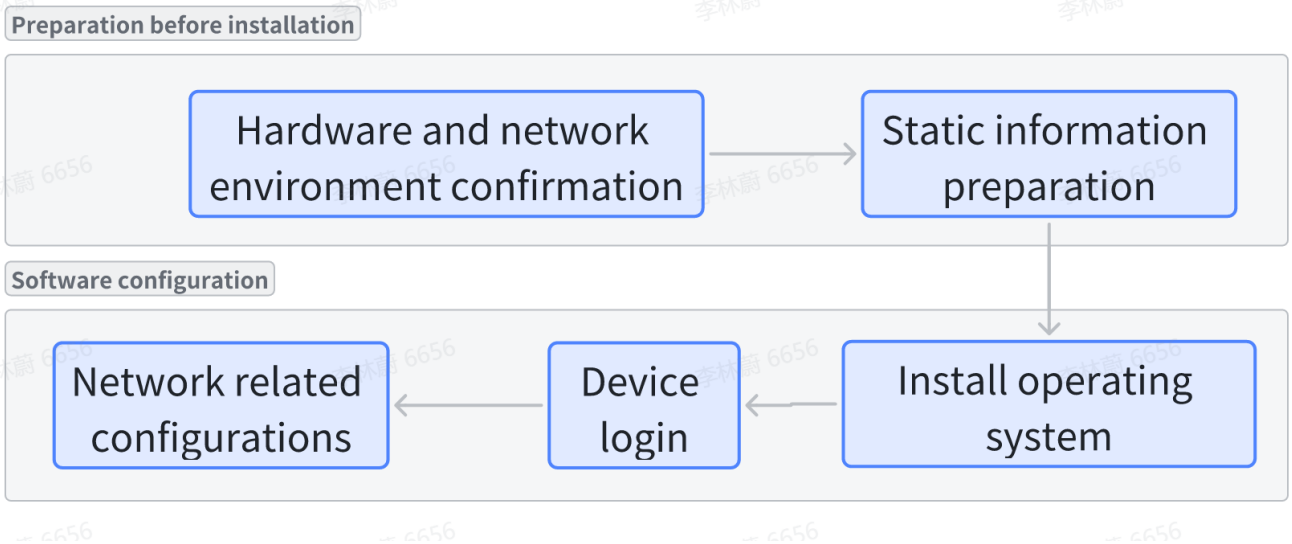

Figure 3-1 Installation Flow Char

The overall installation process described in this manual is as follows: first, prepare for installation, then proceed with software configuration. Once the configuration is complete, the device can be used normally. After that, DPDK and VPP runtime tests can be conducted.

3.3 Software Configuration Checklist

Software Package Name	Content description	Download Location
ET2500-Debian12-Vx.yRz-date.tar	Debian12 OS for ET2500	ET2500-Debian12-Vx.yRz-date.tar
ET2500-Ubuntu24.04-Vx.yRz-date.tar	Ubuntu20.04 OS for ET2500	ET2500-Ubuntu24.04-Vx.yRz-date.tar

Copyright ©2026 Asterfusion. All rights reserved.

ET2500-dpdk-24.03	DPDK for ET2500	ET2500-dpdk-24.03
ET2500-vpp-24.02	VPP for ET2500	ET2500-vpp-24.02

Table 3-2 ET2500 Software Configuration Checklist

3.4 Compatibility List

OS	Kernel	gcc version
Debian12	6.1.67	12.2.0
Ubuntu24.04		

Table 3-3 Compatibility List

Copyright ©2026 Asterfusion. All rights reserved.

4 Preparation before Installation

4.1 Confirmation of Hardware and Network Environment

Serial Number	Description
1	The ET2500 hardware has been completed, including the memory, power supply, fans, and other components.
2	The ET2500 is now ready for power-up.
3	The ET2500 is connected to the terminal host and can be accessed via the serial port.

Table 4-1 Confirmation of Hardware and Network Environment

First, check the ET2500 hardware appearance for any damage, including the memory, power supply, fans, heat sinks, and other components. Then, insert the power cable into the device and turn on the power socket switch to power on the device. Next, prepare an RJ45 to USB serial cable; connect the RJ45 end to the serial port on the right side of the device's front panel and the USB end to the terminal host.

4.2 Preparation of Static Information

Content	Source	Explanation
MAC address to be assigned	Uniformly distributed from the factory or modified on-site	Uniformly distributed from the factory or modified on site
IP address of the device to be deployed	Factory default assignment 192.168.2.100 or modified on-site	For remote access and software updates, use the default IP address if the device is not online

Table 4-2 Preparation of Static Information

Copyright ©2026 Asterfusion. All rights reserved.

5 Software Configuration of ET2500

The ET2500 network appliance is pre-installed with the Debian 12 operating system. For updating or reinstalling the operating system, please refer to Appendix A.

5.1 ET2500 Login

There are two login modes for ET2500: serial port, management network port.

5.1.1 Serial Port Login

On a Windows system, the user can use a serial connection-supported application, such as Mobaxterm or HyperTerminal. Set the connection speed to 115200, and follow the rest of the settings as shown in Figure 5-2, Mobaxterm Advanced Settings, to successfully establish the connection.

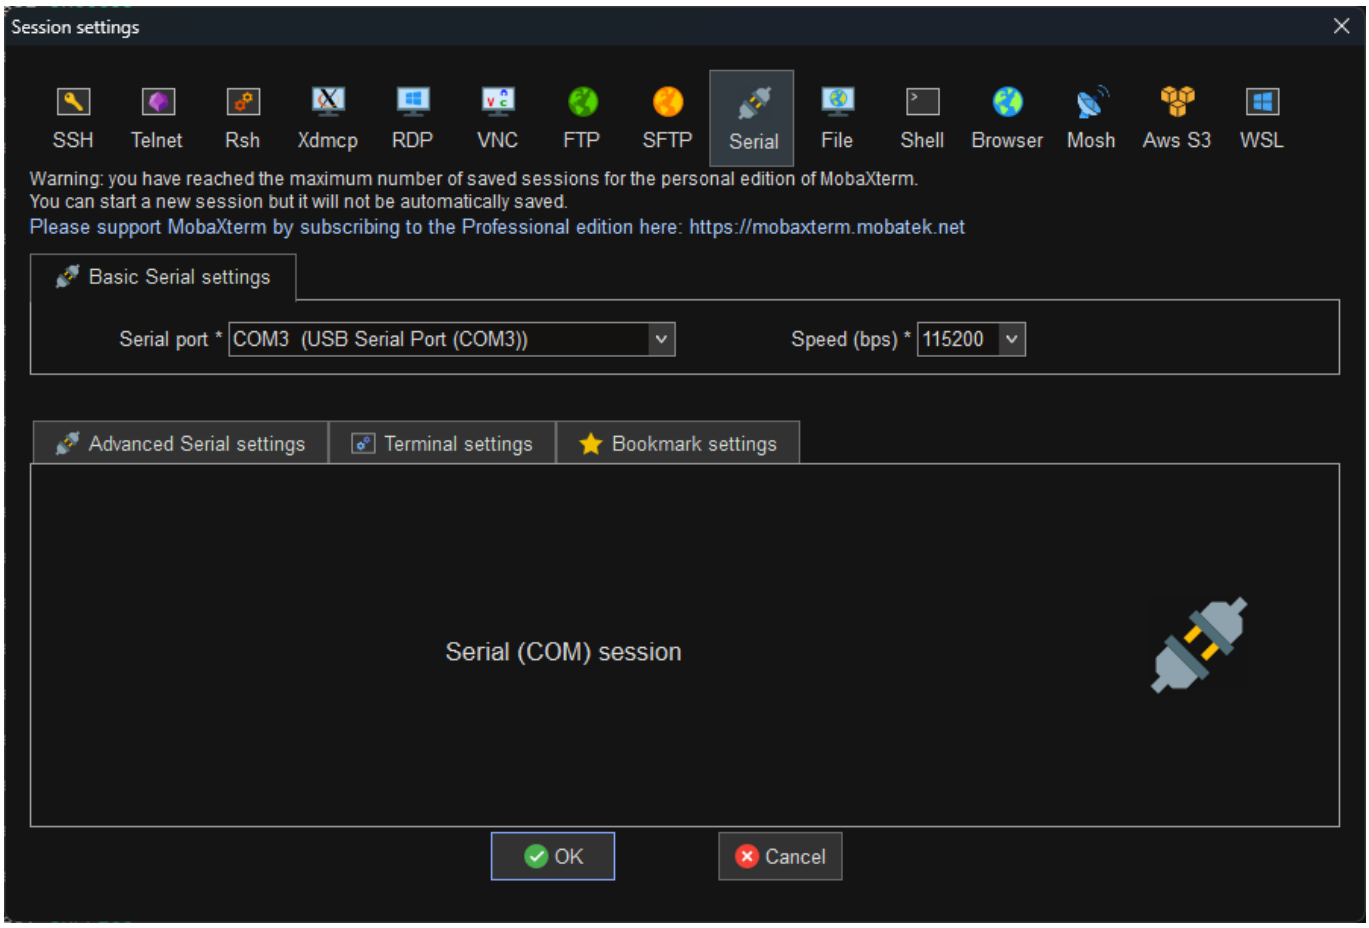

Figure 5-1 Mobaxterm Serial Port Connection

Copyright ©2026 Asterfusion. All rights reserved.

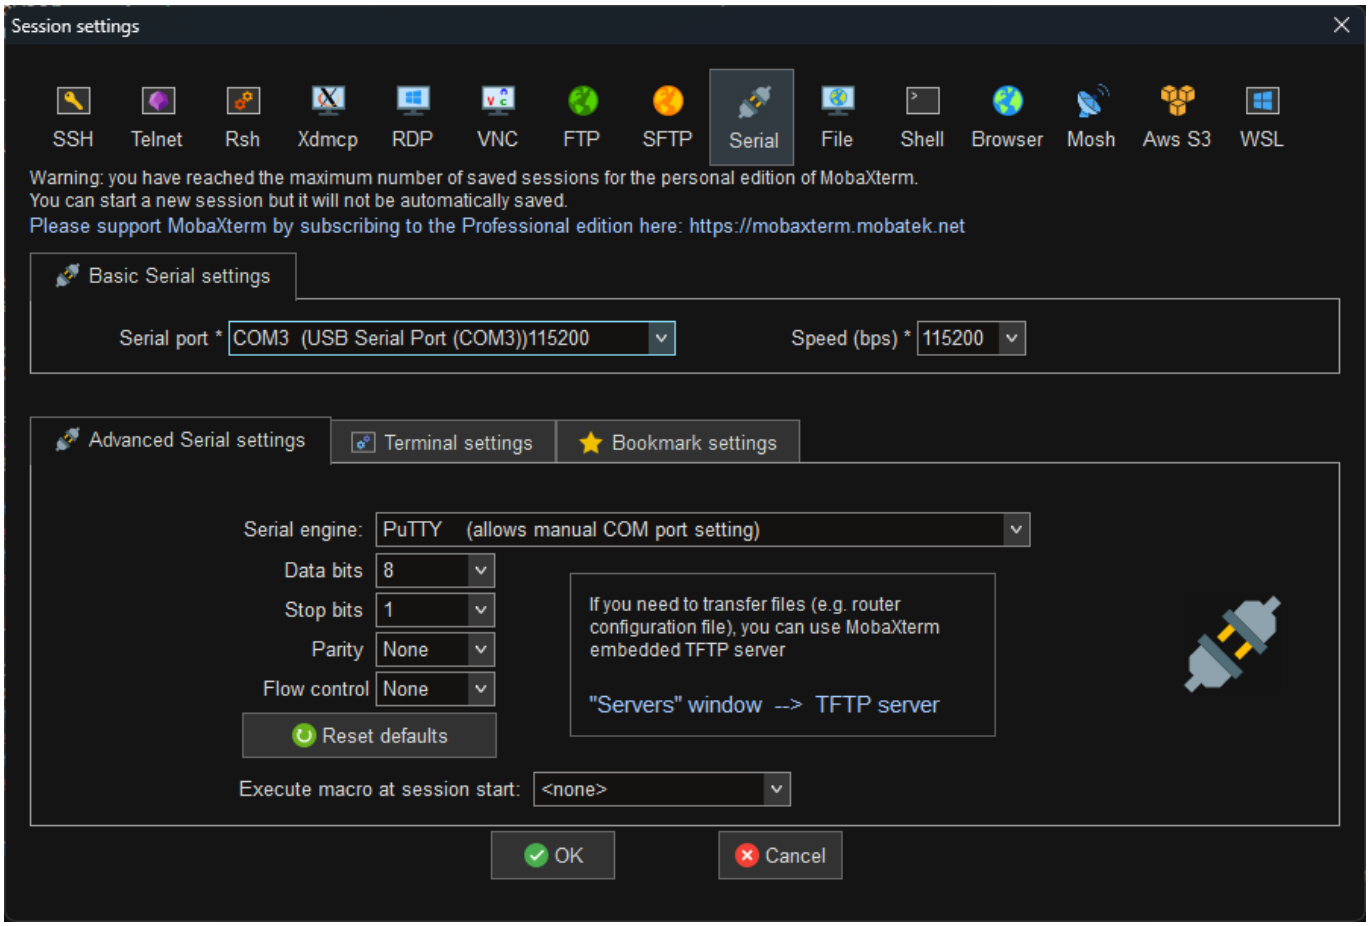

5-2, Mobaxterm Advanced Settings

The configuration for Minicom in Linux is as follows:

Figure 5-3 Configuration for Minicom

The system's default username is **root**, and the password is **passok**.

5.1.2 Management Network Port Login

The default IP configuration for the management port Ethernet0 is **192.168.2.100,** and you can log in to this address via SSH.

You can also configure the IP address manually by logging into the device through the serial port to set

Copyright ©2026 Asterfusion. All rights reserved.

the management port IP address and configure SSH access permissions.

(For detailed network configuration, see section 5.3, "Network-Related Configuration.")

5.2 PCI Address of ET2500

The device's front panel ports correspond to the following PCI addresses:

Front Panel Ethernet Port Numbering	Linux System Network Interface Names	PCI address	Port Type
1	Ethernet0	0002:05:00.0	1GE(RJ45)
2	Ethernet1	0002:02:00.0	1GE(RJ45)
3	Ethernet2	0002:04:00.0	1GE(RJ45)
4	Ethernet3	0002:03:00.0	1GE(RJ45)
5	Ethernet4	0002:09:00.0	1GE(RJ45)
6	Ethernet5	0002:06:00.0	1GE(RJ45)
7	Ethernet6	0002:08:00.0	1GE(RJ45)
8	Ethernet7	0002:07:00.0	1GE(RJ45)
9	Ethernet8	0002:0d:00.0	2.5GE(RJ45)
10	Ethernet9	0002:0a:00.0	2.5GE(RJ45)
11	Ethernet10	0002:0c:00.0	2.5GE(RJ45)
12	Ethernet11	0002:0b:00.0	2.5GE(RJ45)
13	Ethernet12	0002:11:00.0	10GE(SFP+)
14	Ethernet13	0002:10:00.0	10GE(SFP+)
15	Ethernet14	0002:0f:00.0	10GE(SFP+)
16	Ethernet15	0002:0e:00.0	10GE(SFP+)

Table 5-1 Device PCI Address Mapping

5.3 Network-Related Configuration

5.3.1 Static IP Configuration

The default management port is Ethernet0, with a default IP of 192.168.2.100. If you need to modify the IP, follow these steps:

Example: Modify the default IP to 192.168.0.212

Step 1: Check the current system version

```
# cat /etc/os-release
```

If the current system is Debian 12:

Copyright ©2026 Asterfusion. All rights reserved.

Step 2: Modify the /etc/network/interfaces file

Open the file and input the following content, then save and close it:

```
# vim /etc/network/interfaces
```

```
auto Ethernet0
iface Ethernet0
inet static address
192.168.0.212
netmask 255.255.255.0
gateway 192.168.0.1
broadcast 192.168.0.255
```

Step 3: Restart the network service

```
# sudo service networking restart
```

If the current system is Debian 24:

Step 2: Modify the /etc/netplan/50-cloud-init.yaml file

```
# vim /etc/netplan/50-cloud-init.yaml
```

```
network:
   ethernets:
       Ethernet0:
           dhcp4: no
           dhcp6: no
           addresses:
               -
192.168.0.212/24
           routes:
               - to: default
                 via:
192.168.200.1
version: 2
```

Note: Pay attention to indentation when modifying the file.

Step 3: Restart the network service

Copyright ©2026 Asterfusion. All rights reserved.

```
# sudo netplan apply
```

5.3.2 Port Rate Settings for Panel Ports

The device can modify the panel port speed using the ethtool tool:

Panel Port Rate	ethtool Arguement
10M	0x2
100M	0x8
1G	0x20000000000
10G	0x80000000000

Table 5-2 Ethtool Argument for Port Rate Settings

For example, change the panel port Ethernet15:

Modified to 10G:

```
# ethtool -s Ethernet15 advertise 0x80000000000
```

Modified to 1G:

```
# ethtool -s Ethernet15 advertise 0x20000000000
```

Copyright ©2026 Asterfusion. All rights reserved.

6 DPDK Runtime Test

6.1 Compile and Install DPDK

Download the DPDK source code from GitHub. Refer to the following link: 3.3 Software Configuration Checklist

After the download is complete, navigate to the dpdk-24.03 folder:

```
# cd dpdk-24.03
```

Install the required dependencies, such as libnuma-dev:

```
# apt install -y libnuma-dev python3-pyelftools meson
```

Start the compilation process:

```
# meson build -Dmax_lcores=8
# ninja -C build
```

The compiled DPDK program is located in the build directory. If you want to install DPDK to the system, execute:

```
# ninja -C build install
```

6.2 Configuring System HugePage

DPDK uses Huge Pages to improve memory management efficiency and reduce memory access latency, enhancing packet processing performance. Here are the configuration methods for Huge Pages:

The default Huge Page size for the ET2500 is 524288 kB. To allocate 8G of Huge Pages, you need to change the number of Huge Pages to 16:

```
# sysctl -w vm.nr_hugepages=16
```

The command to check the current system's Huge Pages is:

```
# grep Huge /proc/meminfo
```

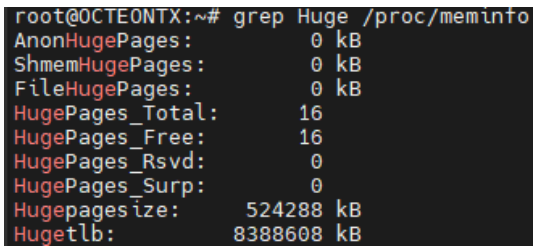

Figure 6-1 check the current system's Huge Pages

Copyright ©2026 Asterfusion. All rights reserved.

Here are the meanings of each field:

HugePages_Total: The total number of Huge Pages configured in the system.

HugePages_Free: The number of Huge Pages that are currently free and available for allocation.

HugePages_Rsvd: The number of Huge Pages that are reserved but not currently in use.

HugePages_Surp: The number of Huge Pages that are allocated but not used by any process.

Hugepagesize: The size of each Huge Page (default 524288kB).

Hugetlb: The total Huge Page memory size.

6.3 Switching Port Drivers

Before running DPDK, it is necessary to switch the relevant port driver from rvu_nicpf to vfio-pci, allowing DPDK to access the port devices directly without going through the operating system. To switch the driver, use the dpdk-24.03/usertools/dpdk-devbind.py script. Here are the operation methods (the following commands should be executed after DPDK is installed; if not installed, include the path to dpdk-devbind.py):

To query the current system for all PCI devices and their status, use the following commands:

```
# dpdk-devbind.py -s
# dpdk-devbind.py –status
```

To bind a device to the DPDK driver, use the following commands:

```
# dpdk-devbind.py -bind vfio-pci 0002:02:00.0
# dpdk-devbind.py --bind=vfio-pci 0002:02:00.0
```

To unbind a device, you can check the help information with these commands:

```
# dpdk-devbind.py -u 0002:02:00.0
# dpdk-devbind.py --unbind 0002:02:00.0
```

To view the help documentation for dpdk-devbind.py, use the following command:

```
# dpdk-devbind.py -h
# dpdk-devbind.py –help
```

6.4 Run Testpmd

To run testpmd using 7 cores with a packet length of 1500 bytes, use the following command:

```
# dpdk-testpmd -l 0-7 -- --nb-cores=7 --txpkts=1500 -i
```

To start packet transmission, enter the following command in the testpmd prompt:

Copyright ©2026 Asterfusion. All rights reserved.

```
testpmd> start tx_first
```

To view information about all ports, use:

```
testpmd> show port info all
```

To see the statistics for all ports, use:

```
testpmd> show port stats all
```

To view the forwarding configuration, enter:

```
testpmd> show config fwd
```

To stop testpmd, use:

```
testpmd> stop
```

To quit testpmd prompt, use:

```
testpmd> quit
```

6.5 DPDK-testpmd Running Example

In this example, DPDK is installed on the system, using the network ports Ethernet14 and Ethernet15 as test ports. Connect Ethernet14 to Ethernet15 using optical fiber or a network cable, and then bind the ports using the DPDK driver.

To simplify the test, the "send then forward" mode (tx_first) will be used.

The network topology for the test is as follows:

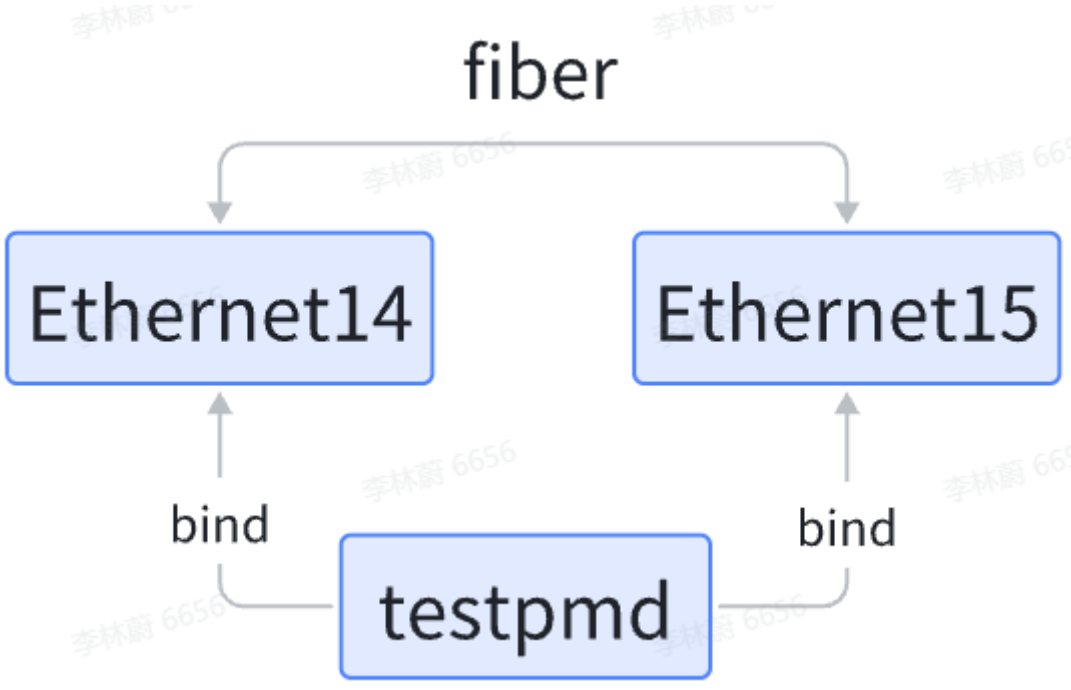

Figure 6-2 Test Network Topology

To begin, compile and install DPDK as described in Section 6.1: Compiling and Installing DPDK. Below are

Copyright ©2026 Asterfusion. All rights reserved.

partial screenshots from the compilation and installation process.

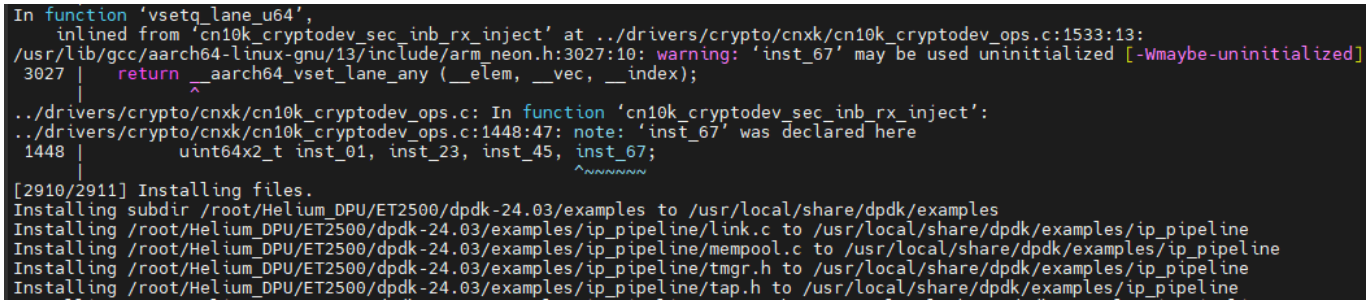

Figure 6-2 Compilation and Installation Process Screenshots

In this example, the command `# ninja -C build install` is executed to install the compiled DPDK into the system, making it easier to run. As shown in the figure, the top part displays the output of the compilation process, while the bottom part shows the output of the installation process.

After that, configure HugePages and drivers, allocate 8GB of HugePages, switch to the DPDK driver, and connect the two test ports using a network cable or optical fiber.

```
# sysctl -w vm.nr_hugepages=16
# grep Huge /proc/meminfo
```

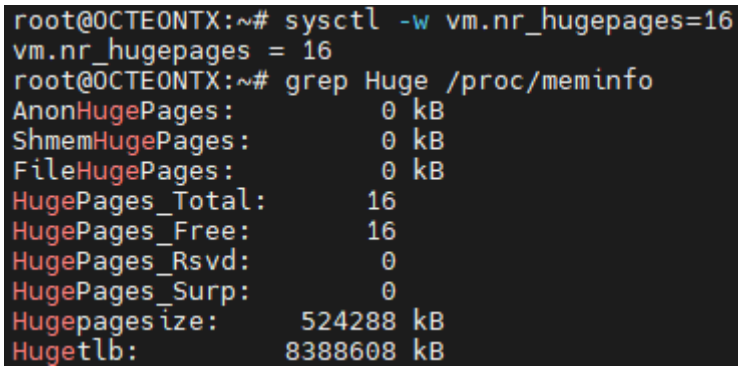

Figure 6-3 Allocating 16 HugePages

```
# dpdk-devbind.py -b vfio-pci 0002:0e:00.0 0002:0f:00.0
# dpdk-devbind.py -s
```

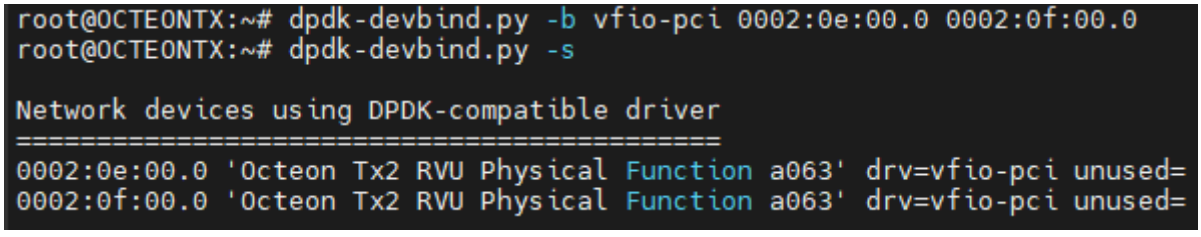

Figure 6-4 Switching Port Drivers

Then run testpmd:

```
# dpdk-testpmd -l 0-7 -- --nb-cores=7 --txpkts=1500 -i
```

Copyright ©2026 Asterfusion. All rights reserved.

```
root@OCTEONTX:~# dpdk-testpmd -l 0-7 -- --nb-cores=7 --txpkts=1500 -i
EAL: Detected CPU lcores: 8
EAL: Detected NUMA nodes: 1
EAL: Detected static linkage of DPDK
EAL: Multi-process socket /var/run/dpdk/rte/mp_socket
EAL: Selected IOVA mode 'VA'
EAL: VFIO support initialized
EAL: Using IOMMU type 1 (Type 1)
EAL: Probe PCI driver: net_cn10k (177d:a063) device: 0002:0e:00.0 (socket 0)
CNXK: RoC Model: cn10kb_a0 (HW_PLATFORM)
EAL: Probe PCI driver: net_cn10k (177d:a063) device: 0002:0f:00.0 (socket 0)
CNXK: Skipping device, mcs device already probed
TELEMETRY: No legacy callbacks, legacy socket not created
Interactive-mode selected
testpmd: create a new mbuf pool <mb_pool_0>: n=203456, size=2176, socket=0
testpmd: preferred mempool ops selected: cn10k_mempool_ops
Configuring Port 0 (socket 0)
CNXK: Port 0: Link Up - speed 10000 Mbps - full-duplex

Port 0: link state change event
Port 0: CE:6C:67:30:0C:02
Configuring Port 1 (socket 0)
CNXK: Port 1: Link Up - speed 10000 Mbps - full-duplex

Port 1: link state change event
Port 1: 0E:9F:22:40:A5:C1
Checking link statuses...
Done
testpmd>
```

Figure 6-5 run testpmd

As shown in the figure, the startup was successful, and commands can be entered normally in the testpmd prompt:

```
testpmd> start tx_first
```

```
testpmd> start tx_first
io packet forwarding - ports=2 - cores=2 - streams=2 - NUMA support enabled, MP allocation mode: native
Logical Core 1 (socket 0) forwards packets on 1 streams:
  RX P=0/Q=0 (socket 0) -> TX P=1/Q=0 (socket 0) peer=02:00:00:00:00:01
Logical Core 2 (socket 0) forwards packets on 1 streams:
  RX P=1/Q=0 (socket 0) -> TX P=0/Q=0 (socket 0) peer=02:00:00:00:00:00

  io packet forwarding packets/burst=32
  nb forwarding cores=7 - nb forwarding ports=2
  port 0: RX queue number: 1 Tx queue number: 1
    Rx offloads=0x0 Tx offloads=0x10000
    RX queue: 0
      RX desc=4096 - RX free threshold=0
      RX threshold registers: pthresh=0 hthresh=0  wthresh=0
      RX Offloads=0x0
    TX queue: 0
      TX desc=512 - TX free threshold=0
      TX threshold registers: pthresh=0 hthresh=0  wthresh=0
      TX offloads=0x10000 - TX RS bit threshold=0
```

Figure 6-6 Start Forwarding

When forwarding is started, display the port statistics:

```
testpmd> show port stats all
```

Copyright ©2026 Asterfusion. All rights reserved.

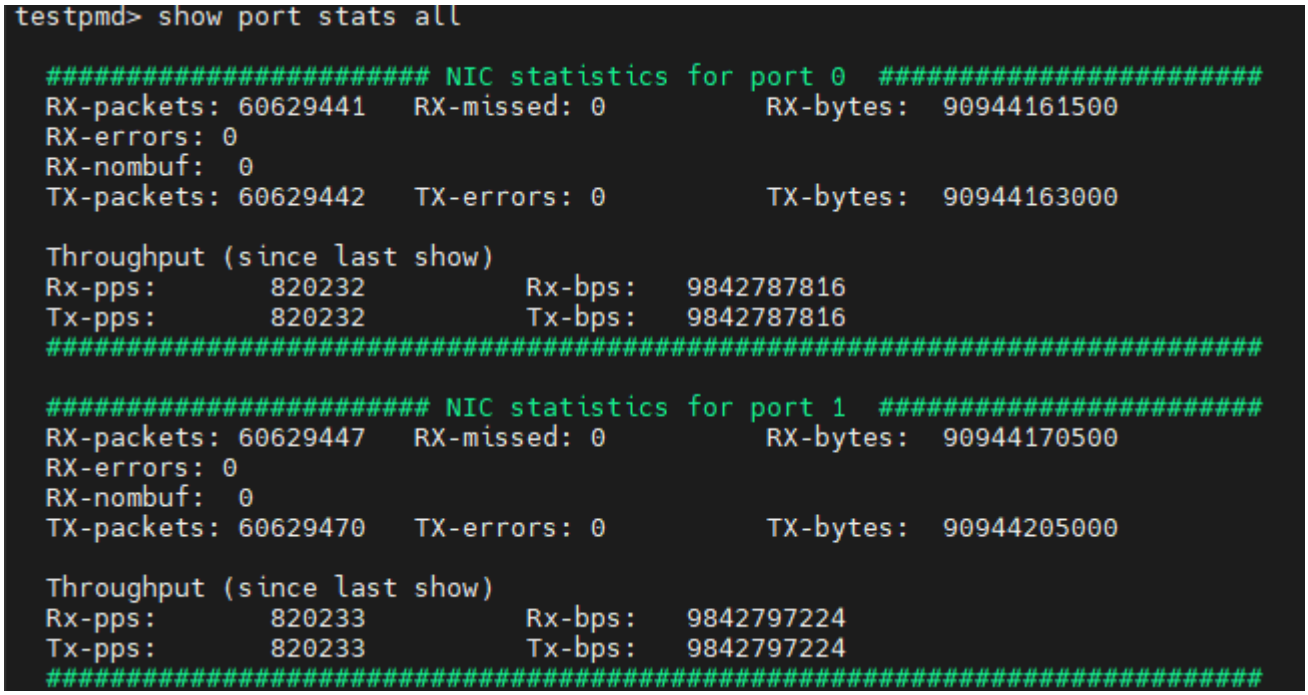

Figure 6-7 display the port statistics

Use this command to view the packet counts, transmission rates, and packet loss for all ports. As shown in the figure, the current port rate is 9.8 Gb/s. Finally, use the stop command to stop forwarding and the quit command to exit the testpmd command line.

```
testpmd> stop
```

```
testpmd> quit
```

Copyright ©2026 Asterfusion. All rights reserved.

7 VPP Runtime Test

7.1 Compile and Install VPP

Download the VPP source code from GitHub, referring to the address in section 3.3 Software Configuration Checklist.

Once the download is complete, navigate to the vpp-24.02 folder and execute the build script:

```
# cd vpp-24.02
# ./vpp-build.sh -release
```

During the compilation process, you may encounter some options that require manual confirmation. Follow the prompts and input y or take action based on the specific situation.

```
make[1]: Leaving directory '/root/Helium_DPU/ET2500/vpp-24.02/build-root'
All operations completed successfully.
root@OCTEONTX:~/Helium_DPU/ET2500/vpp-24.02#
```

Figure 7-1 Compile VPP

Printing this message indicates successful compilation.

Install VPP:

```
# make pkg-deb
# dpkg -i build-root/*.deb
```

7.2 Configuring System HugePage

Refer to section 6.2 Configuring System HugePage.

7.3 Switching Port Drivers

The script path for switching drivers is:

`vpp-24.02/extras/vpp_config/scripts/dpdk-devbind.py`

Refer to section 6.3 Switching Port Drivers.

7.4 Edit Configuration File

The configuration file can refer to the example configuration file provided by VPP:

`vpp-24.02/build-root/install-vpp-native/vpp/etc/vpp/onp-startup.conf`

Note: This configuration file is generated only after the compilation is complete.

Copyright ©2026 Asterfusion. All rights reserved.

The following are explanations for some fields in the configuration file:

unix Configuration Block: Configures some basic behaviors of VPP.

api-trace Configuration Block: Controls the tracing functionality of API calls. This feature helps with debugging and tracking API calls.

cpu Configuration Block: Configures the CPU cores used by VPP.

- **main-core Block**: Sets the CPU core to which the main thread is bound, defaulting to core 1.
- **corelist-workers Block**: Sets the range of CPU cores to which worker threads are bound.

buffers Configuration Block: Adjusts buffer allocation settings.

onp Configuration Block: Configures whitelist devices and queue parameters.

dev: Specifies the PCIe address of the device.

plugins Configuration Block: Enables or disables plugins within VPP.

statseg Configuration Block: Configures the statistics segment for monitoring and collecting statistical data.

7.5 Run VPP

After writing the configuration file, use the -C parameter to specify the configuration file for VPP and run VPP:

```
# vpp -C start_ET2500.conf
```

For additional commands, please refer to the VPP documentation:https://s3-docs.fd.io/vpp/24.02/

7.6 VPP Running Example

This run example uses another device, TRex, to send packets, while VPP enables Layer 3 routing forwarding. Ethernet12 and Ethernet13 are bound to VPP and connected to another device.

The test network topology is as follows:

Copyright ©2026 Asterfusion. All rights reserved.

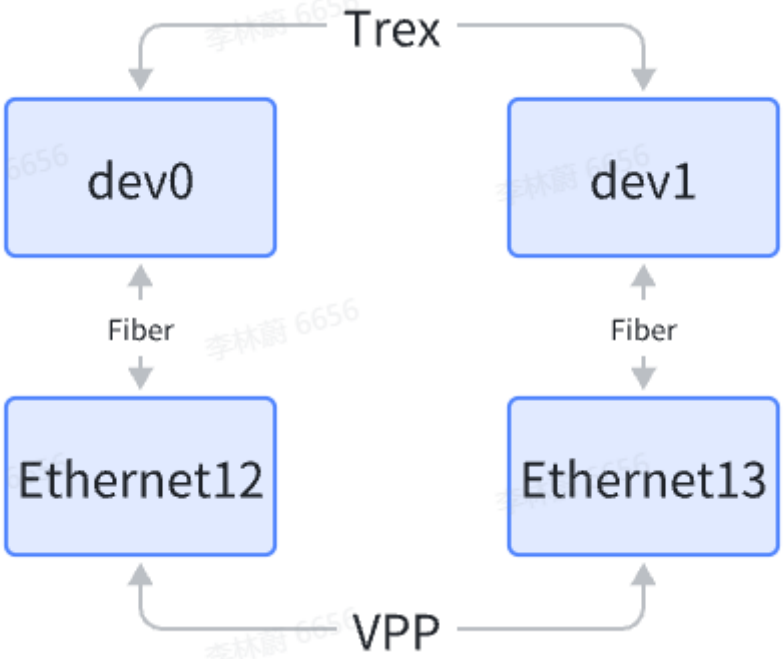

Figure 7-2 Test Network Topology

The VPP configuration file is as follows

```
unix {
  interactive
}

api-trace {
  on
}

memory {
  main-heap-size 1G
  main-heap-page-size default-hugepage
}

cpu {
  main-core 1
  corelist-workers 2-3
```

Copyright ©2026 Asterfusion. All rights reserved.

```
}

buffers {
  buffers-per-numa 128000
  page-size default-hugepage
}

plugins {
  plugin dpdk_plugin.so { disable }
  plugin onp_plugin.so { enable }
  plugin ikev2_plugin.so { enable }
  plugin idpf_plugin.so { enable }
  plugin oddbuf_plugin.so { enable }
}

onp {
  dev 0002:11:00.0 {
    name Ethernet12
  }
  dev 0002:10:00.0 {
    name Ethernet13
  }
}
statseg {
  size 1G
}
```

:

Start VPP and Configure:

```
# ./vpp -c ../ET2500_startup.conf
vppctl> set interface ip address Ethernet12 10.10.14.2/24
vppctl> set interface ip address Ethernet13 10.10.15.2/24
```

Copyright ©2026 Asterfusion. All rights reserved.

```
vppctl> set interface state Ethernet12 up
vppctl> set interface state Ethernet13 up
vppctl> set interface mac address Ethernet12 aa:bb:cc:dd:ee:13
vppctl> set interface mac address Ethernet13 aa:bb:cc:dd:ee:14
vppctl> set ip neighbor Ethernet12 10.10.13.1 d2:d2:1c:43:a6:13
vppctl> set ip neighbor Ethernet13 10.10.14.1 d2:d2:1c:43:a6:14
```

At this point, start the packet generation software Trex and check the port status:

```
vppctl> show interface
```

Copyright ©2026 Asterfusion. All rights reserved.

8 AI Module User Guide

8.1 AI Module Overview

This chapter applies only to configurations with the optional AI module. If your platform does not include this module, you can safely skip this chapter.

The AI Card M.2 is a high-compute module designed for large-model inference at the edge. It offers a compact form factor, high performance, and low power consumption. It can deliver up to 160 TOPS of elastic compute power and 100 TFLOPS @ bFP16 computing capability. The product is equipped with 24 GB of LPDDR5/LPDDR5X memory and uses an M.2 interface design that supports ET2500, ET2700, and ET3600 devices.

The AI Card M.2 has already been adapted to mainstream large models (such as Qwen, GPT, and Gemma) and is suitable for building AI-capable edge network devices.

Software and Driver List

Package Name	Description
docker-ai-agent.tar	Provides llama-server and Hermes service
xh2a_drv.ko	AI Card driver, installed by default in the factory OS

8.2 Installation and Configuration

8.2.1 Check Driver and Device

Run the following commands in sequence to verify device status.

****Check whether the driver is loaded:****

ƀash

```
lsmod | grep xh2a_drv
```

`

Expected output:

ƀash

```
xh2a_drv            327680  6
```

`

****Check whether the PCIe device is recognized:****

ƀash

Copyright ©2026 Asterfusion. All rights reserved.

```
lspci -vvv -d 1f6b:0c00
```

Expected output (key fields):

```
0003:01:00.0 Memory controller: Device 1f6b:0c00
      Kernel modules: xh2a_drv
```

8.2.2 Deploy the docker-ai-agent Image

Download model files

-Download the .gguf model file from the cloud disk: [download models](https://drive.cloudswitch.io/external/b53840a1353fce8076a6f5646e0c023447336d011c3a23653c1e6fca52e88cc2)

- Place the file in /opt/models

Create a symbolic link

```
mkdir -p /opt/models
ln -s /opt/models/qwen3.5_35b.gguf /opt/models/model.gguf
```

Import the Docker image

-Download the docker-ai-agent image: [download ai-agent](https://drive.cloudswitch.io/external/4b1d6498c08b1e3f221abbb6a67a0649388f0dfe2f2535738844108db25d3612/download)

```
docker image load -i docker-ai-agent.tar
```

Create and start the container

```
docker run -it --privileged -v /opt/models:/opt/models --name <container_name>
docker-ai-agent:latest /bin/bash
```

Copyright ©2026 Asterfusion. All rights reserved.

8.2.3 Check Device Version and Serial Number

After entering the container, run:

ðash

```
hm_smi -a
```

`

Expected output (simplified):

ðash

```
HMSW_Version   : V1.2.0
Driver_Version : V1.2.0
Model          : LQ50-24GB
Firmware_Version : V1.1.1
SN             : 0104020100002026000400000350
DDR_Memory_Total : 24448.0MB
Core_Num       : 2
Cur_Ipu_Freq   : 1400.0 Mhz
```

8.2.4 Start the llama-server Service

****Start the service****

ðash

```
supervisorctl start llama-server
```

`

****Check service status****

ðash

```
supervisorctl status
```

`

Expected status: `llama-server RUNNING`

****Verify port listening****

ðash

Copyright ©2026 Asterfusion. All rights reserved.

```
netstat -anpt | grep 8080
```

8.2.5 Test API Functionality

Replace the IP address with the management IP address you configured:

bash

```
curl --request POST \
  --url http://192.168.1.182:8080/v1/chat/completions \
  --header "Content-Type: application/json" \
  --data  '{"model":  "qwen3.5-30b",  "messages":  [{"role":  "user",  "content":
"Test"}], "stream": false}'
```

Expected response includes (simplified):

json

```
{
  "choices": [
    {
      "finish_reason": "stop",
      "message": {
        "role": "assistant",
        "content": "Test received!"
      }
    }
  ],
  "usage": {
    "completion_tokens": 184,
    "prompt_tokens": 11,
    "total_tokens": 195
  }
}
```

Copyright ©2026 Asterfusion. All rights reserved.

8.2.6 Configure Hermes

Refer to the official documentation: [Hermes setup](https://hermes-agent.app/en/quickstart)

****Configure the API address****

Replace the IP address with the management IP address you configured:

`text

provider: custom

base_url: http://192.168.1.182:8080/v1

`

****Restart Hermes****

`bash

supervisorctl restart hermes

`

Copyright ©2026 Asterfusion. All rights reserved.

8.3 Troubleshooting

Issue	Solution
`lsmod` returns no output	Check whether the driver is loaded; if not loaded, use `modprobe xh2a_drv`
`curl` connection refused	Confirm that `llama-server` is running and listening on port 8080
`docker run` permission error	Ensure you are running as root or a sudo user, or add Docker group permissions
Out of memory error	Check available memory displayed by `hm_smi -a`

SMI Tool Description

The SMI (System Management Interface) tool is used to manage and monitor the product. Users can obtain real-time device information such as power, temperature, memory, and IPU cores through the command line, and perform DVFS (Dynamic Voltage and Frequency Scaling) configuration, device reset, and other management operations.

8.3.1 Usage Instructions

After the driver is installed, the SMI tool can be used from any path to get device information. The basic command format is as follows:

`bash`

```
hm_smi [OPTION1 [ARG1]] [OPTION2 [ARG2]] ...
```

8.3.2 Command Parameter Description

| Parameter | Description |

| ---- | ---- |

| `-a, --all` | View information for all current logical devices |

| `-d <device_id>, --device <device_id>` | View information for the logical device with ID `<device_id>` |

| `-f <FIELD> [<FIELD> ...], --fields` | View specified fields of device information (supports multiple fields separated by spaces; must be used together with `-a`) |

| `-g [{ondemand, performance, powerlimit}], --gov` | Set the device DVFS mode |

| `-l <sec> [options], --loop` | Print device information in a loop at the specified interval |

| `-lc LOCK_CLOCKS, --lock_clocks` | Set the IPU core clock frequency range |

| `-r <device_id>, --reset` | Reset the IPU core for the specified logical device |

Copyright ©2026 Asterfusion. All rights reserved.

| -rd <device_id>, --reset_device | Reset the logical device with ID <device_id> (Note: LM5070 and LM5050 do not support this reset method; host reboot is required) |

| -pl POWERLIMIT, --powerlimit | Set the power range for the AI processor card |

| -o <file_name>, --output | Save command output to a file (supports .txt, .json, .xml formats) |

| -v, -vv, -vvv | Set log level (-v basic, -vv detailed, -vvv debug) |

| -h, --help | View command help information |

8.3.3 Output Information Description (Key Fields)

| Field | Description |

| ---- | ---- |

| Build_Time | SMI tool build time |

| HMSW_Version / HM_SMI_Version | Software platform version and tool version |

| Driver_Version / Vendor | Host driver version and device vendor information |

| BDF / Dev | PCIe BDF (Bus, Device, Function) information and logical ID |

| Power_Management | Power management status, including DVFS Mode, Cur_Ipu_Freq, Max_Power_Limit, etc. |

| DDR_Memory_Infos | Current logical device memory information, including free and total capacity |

| Temperature | Real-time temperature of each IPU core and DDR module (℃) |

| Board_Power | Real-time average power consumption of the AI processor card (W) |

| IPU_Infos | Includes core count, frequency, voltage, and average utilization |

8.3.4 Common Query Examples

- View information for all devices:

bash

```
hm_smi -a
```

`

- View information for a specified logical device:

bash

```
hm_smi -d <device_id>
```

`

Copyright ©2026 Asterfusion. All rights reserved.

- View specified field information:

ðash

```
hm_smi -a -f Driver_Version
```

```
hm_smi -a -f Cur_BandWidth DDR_Memory_Free -d 0
```

`

- Loop display device information every 5 seconds:

ðash

```
hm_smi -l 5 -d 0
```

`

8.3.5 Advanced Settings and Features

****Set device DVFS mode****

Supported modes:

- þerformance:` default mode, maximum frequency

- òndemand:` dynamically adjust as needed

- þowerlimit:` adjust according to power limit

Example:

ðash

```
hm_smi -g powerlimit -pl 4w -l 30 -a
```

`

****Set IPU core frequency range****

ðash

```
hm_smi -lc 1000 700 -a
```

`

****Reset a specified IPU core****

ðash

```
hm_smi -r 0
```

`

****Reset the entire device (requires root permissions)****

ðash

```
sudo hm_smi -rd 0
```

Copyright ©2026 Asterfusion. All rights reserved.

`

Save output to a file

bash

```
hm_smi -o smi_return.txt
```

`

Set log level (output the most detailed debug information)

bash

```
hm_smi -vvv
```

`

9 Appendix A Operating System Installation

9.1 Boot Installation from USB

9.1.1 Preparation Work

Connect the device via the serial port.

Backup important files and data from the eMMC device.

Prepare a USB storage device containing a complete Linux file system and Linux kernel image, such as the BusyBox system.

Download the latest Linux image compressed package to the prepared storage device, for example ET2500-Debian12-V1.0R2-20240903.tar.

9.1.2 Boot Parameter Settings

In this step, the boot disk needs to be changed from eMMC to a USB device.

First, restart the device:

```
# reboot
```

When the message "Autoboot in 5 seconds" appears, press the **S** and **T** keys on the keyboard to enter the Uboot interface.

In the Uboot interface, print the current boot parameters:

```
et2500> printenv
```

Insert the prepared USB and start Uboot's USB:

Copyright ©2026 Asterfusion. All rights reserved.

```
et2500> usb start
```

If you see the message "0 Storage Device(s) found," you need to run:

```
et2500> usb reset
```

Repeat until the storage device appears.

If you are unsure about the device, you can check the USB storage contents using the ls command. For example, if the current storage device is formatted as ext4, you can use ext4ls:

```
et2500> usb start
et2500> ext4ls usb 0:1
```

This represents the first partition of the 0th USB device.

When there are multiple storage devices, they are sorted in order as **sda**, **sdb**, **sdc**, etc.

When a storage device has multiple partitions, they are sorted as **sda1**, **sda2**, **sda3**, etc.

Modify the bootargs to specify the prepared storage device, for example root=/dev/sda1:

```
et2500> editenv bootargs
```

The Uboot command line will display the current bootargs, modify it to:

```
console=ttyAMA0,115200n8 earlycon=pl011,0x87e028000000 maxcpus=8 rootwait rw
root=/dev/sda1 coherent_pool=16M net.ifnames=0
```

Note: After modifying with editenv bootargs, use `printenv` to verify the changes.

(Optional) Store the environment variables: You can save the environment variables to non-volatile storage:

```
et2500> saveenv
```

9.1.3 Boot from USB Device

Load the kernel onto the device:

```
et2500> ext4load usb 0:1 $loadaddr boot/ET2500-uboot.img
```

This represents the ET2500-uboot.img kernel loacted in the /boot folder of the first partition of the 0th USB device. Adjust according to your situation.

Start BusyBox:

```
et2500> booti $loadaddr - $fdtcontroladdr
```

9.1.4 Install Linux

Extract the installation package into a new folder:

```
# mkdir temp
```

Copyright ©2026 Asterfusion. All rights reserved.

```
# tar -xvf ET2500-Debian12-V1.0R2-20240903.tar -C temp
# cd temp
```

Note: If you are using a single USB device to install multiple systems, you need to extract each system into an empty folder each time.

After successful extraction, you should see the following folders:

```
driver.tar.gz
eMMCAutoInstall.sh
ET2500-uboot.img
extra_files
Image
partition.sh
rootfs_debian12.tar.gz
rootfs_patch.sh
```

Run the eMMCAutoInstall.sh script:

```
# ./eMMAutoInstall.sh
```

If you do not have execution permissions, you need to run:

```
# sudo chmod +x eMMAutoInstall.sh
```

Wait for the installation to complete, then restart the device to boot into the newly installed system.

9.2 Boot Installation from Network

9.2.1 Preparation Work

Connect to the device via the serial port.

Backup important files and data from the eMMC device.

Prepare a TFTP server and connect it to the device's management network port.

Prepare the BusyBox system and the Linux installation package, such as ET2500-Debian12-V1.0R2-20240903.tar

9.2.2 Download BusyBox and System Installation Package

The default management network port for the ET2500 device is **Ethernet0** (rvu_pf#3). Execute `et2500> ethlist` to confirm whether the rvu_pf#3 port exists. If it does not, enter `et2500> usb reset` to activate the management network port.

Copyright ©2026 Asterfusion. All rights reserved.

Note: After rebooting, you need to re-confirm whether the management network port is activated.

Place the BusyBox image on the TFTP server, configure the management network port's IP and TFTP server IP to ensure both are on the same subnet. Then, use TFTP to download the BusyBox image and enter BusyBox.

```
et2500> usb reset
et2500> setenv ipaddr {ip}
et2500> setenv serverip {tftp_server_ip}
et2500> ethact
et2500> tfptboot $loadaddr ET2500-busybox-Vx.yRz.img;booti $loadaddr -
$fdtcontroladdr
```

Enter the BusyBox terminal to download the Linux system package:

```
# ifconfig Ethernet0 <ip>
# ip route add default via <gateway_ip>
# scp root@<server_ip>:~/ET2500- Debian12-Vx.yRz.tar ./
```

9.2.3 Install Linux

Refer to Section 8.1.4 Install Linux.

Copyright ©2026 Asterfusion. All rights reserved.

10 Appendix B: POE Port Configuration

10.1 POE Configuration Commands

Initialize the POE system; this needs to be executed only once after each startup:

```
# poe_device_control poe init
```

Display the current POE system status:

```
# poe_device_control poe show
```

Enable POE on the specified port. Use "all" to enable Power over Ethernet on all ports.

```
# poe_device_control poe enable|disable all|EthernetX
```

For example:

```
# poe_device_control poe enable all
```

Set the maximum power for the specified port. Valid values include: 60w, 30w, 15w, 30w-at.

```
# poe_device_control maxpower {60w|30w|15w|30w-at} all|EthernetX
```

For example:

```
# poe_device_control maxpower 30w Ethernet4
```

Set the priority for the specified port. Valid values include: low, high, critical.

```
# poe_device_control priority {low|high|critical} all|EthernetX
```

Enable or disable legacy detection on the specified port.

```
# poe_device_control legacydetect {enable|disable} all|EthernetX
```

Control the LED status of the specified port.

```
# poe_device_control led on|off all|EthernetX
```

10.2 Notes

Enabling and Disabling:

Power can only be enabled if the total device power is sufficient.

Power can only be enabled if the current available power is greater than or equal to the required power; otherwise, the port status will indicate 1F (overload).

There is a certain delay when enabling, disabling, or switching device states; wait for a while before running the show command.

The 2.5G port only has a 60w-at setting.

Copyright ©2026 Asterfusion. All rights reserved.