



**DPDK**

**DATA PLANE DEVELOPMENT KIT**

**FAQ**

***Release 18.11.0***

November 27, 2018

## CONTENTS

<b>1</b>	<b>What does “EAL: map_all_hugepages(): open failed: Permission denied Cannot init memory” mean?</b>	<b>2</b>
<b>2</b>	<b>If I want to change the number of hugepages allocated, how do I remove the original pages allocated?</b>	<b>3</b>
<b>3</b>	<b>If I execute “l2fwd -l 0-3 -m 64 -n 3 – -p 3”, I get the following output, indicating that there are no socket 0 hugepages to allocate the mbuf and ring structures to?</b>	<b>4</b>
<b>4</b>	<b>I am running a 32-bit DPDK application on a NUMA system, and sometimes the application initializes fine but cannot allocate memory. Why is that happening?</b>	<b>5</b>
<b>5</b>	<b>On application startup, there is a lot of EAL information printed. Is there any way to reduce this?</b>	<b>6</b>
<b>6</b>	<b>How can I tune my network application to achieve lower latency?</b>	<b>7</b>
<b>7</b>	<b>Without NUMA enabled, my network throughput is low, why?</b>	<b>8</b>
<b>8</b>	<b>I am getting errors about not being able to open files. Why?</b>	<b>9</b>
<b>9</b>	<b>VF driver for IXGBE devices cannot be initialized</b>	<b>10</b>
<b>10</b>	<b>Is it safe to add an entry to the hash table while running?</b>	<b>11</b>
<b>11</b>	<b>What is the purpose of setting iommu=pt?</b>	<b>12</b>
<b>12</b>	<b>When trying to send packets from an application to itself, meaning smac==dmac, using Intel(R) 82599 VF packets are lost.</b>	<b>13</b>
<b>13</b>	<b>Can I split packet RX to use DPDK and have an application’s higher order functions continue using Linux pthread?</b>	<b>14</b>
<b>14</b>	<b>Is it possible to exchange data between DPDK processes and regular userspace processes via some shared memory or IPC mechanism?</b>	<b>15</b>
<b>15</b>	<b>Can the multiple queues in Intel(R) I350 be used with DPDK?</b>	<b>16</b>
<b>16</b>	<b>How can hugepage-backed memory be shared among multiple processes?</b>	<b>17</b>
<b>17</b>	<b>Why can’t my application receive packets on my system with UEFI Secure Boot enabled?</b>	<b>18</b>

This document contains some Frequently Asked Questions that arise when working with DPK.

**WHAT DOES “EAL: MAP\_ALL\_HUGEPAGES(): OPEN FAILED:  
PERMISSION DENIED CANNOT INIT MEMORY” MEAN?**

This is most likely due to the test application not being run with sudo to promote the user to a superuser. Alternatively, applications can also be run as regular user. For more information, please refer to DPDK Getting Started Guide.

## IF I WANT TO CHANGE THE NUMBER OF HUGE PAGES ALLOCATED, HOW DO I REMOVE THE ORIGINAL PAGES ALLOCATED?

The number of pages allocated can be seen by executing the following command:

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

Once all the pages are mapped by an application, they stay that way. If you start a test application with less than the maximum, then you have free pages. When you stop and restart the test application, it looks to see if the pages are available in the `/dev/huge` directory and maps them. If you look in the directory, you will see `n` number of 2M pages files. If you specified 1024, you will see 1024 page files. These are then placed in memory segments to get contiguous memory.

If you need to change the number of pages, it is easier to first remove the pages. The `usertools/dpdk-setup.sh` script provides an option to do this. See the “Quick Start Setup Script” section in the DPDK Getting Started Guide for more information.

**IF I EXECUTE “L2FWD -L 0-3 -M 64 -N 3 – -P 3”, I GET THE FOLLOWING OUTPUT, INDICATING THAT THERE ARE NO SOCKET 0 HUGE PAGES TO ALLOCATE THE MBUF AND RING STRUCTURES TO?**

I have set up a total of 1024 Hugepages (that is, allocated 512 2M pages to each NUMA node).

The -m command line parameter does not guarantee that huge pages will be reserved on specific sockets. Therefore, allocated huge pages may not be on socket 0. To request memory to be reserved on a specific socket, please use the –socket-mem command-line parameter instead of -m.

## **I AM RUNNING A 32-BIT DPDK APPLICATION ON A NUMA SYSTEM, AND SOMETIMES THE APPLICATION INITIALIZES FINE BUT CANNOT ALLOCATE MEMORY. WHY IS THAT HAPPENING?**

32-bit applications have limitations in terms of how much virtual memory is available, hence the number of hugepages they are able to allocate is also limited (1 GB size). If your system has a lot (>1 GB size) of hugepage memory, not all of it will be allocated. Due to hugepages typically being allocated on a local NUMA node, the hugepages allocation the application gets during the initialization depends on which NUMA node it is running on (the EAL does not affinity cores until much later in the initialization process). Sometimes, the Linux OS runs the DPDK application on a core that is located on a different NUMA node from DPDK master core and therefore all the hugepages are allocated on the wrong socket.

To avoid this scenario, either lower the amount of hugepage memory available to 1 GB size (or less), or run the application with taskset affinity to a would-be master core.

For example, if your EAL coremask is 0xff0, the master core will usually be the first core in the coremask (0x10); this is what you have to supply to taskset:

```
taskset 0x10 ./l2fwd -l 4-11 -n 2
```

In this way, the hugepages have a greater chance of being allocated to the correct socket. Additionally, a `--socket-mem` option could be used to ensure the availability of memory for each socket, so that if hugepages were allocated on the wrong socket, the application simply will not start.

**ON APPLICATION STARTUP, THERE IS A LOT OF EAL INFORMATION PRINTED. IS THERE ANY WAY TO REDUCE THIS?**

Yes, the option `--log-level=` accepts either symbolic names (or numbers):

1. emergency
2. alert
3. critical
4. error
5. warning
6. notice
7. info
8. debug



## **HOW CAN I TUNE MY NETWORK APPLICATION TO ACHIEVE LOWER LATENCY?**

Traditionally, there is a trade-off between throughput and latency. An application can be tuned to achieve a high throughput, but the end-to-end latency of an average packet typically increases as a result. Similarly, the application can be tuned to have, on average, a low end-to-end latency at the cost of lower throughput.

To achieve higher throughput, the DPDK attempts to aggregate the cost of processing each packet individually by processing packets in bursts. Using the testpmd application as an example, the “burst” size can be set on the command line to a value of 32 (also the default value). This allows the application to request 32 packets at a time from the PMD. The testpmd application then immediately attempts to transmit all the packets that were received, in this case, all 32 packets. The packets are not transmitted until the tail pointer is updated on the corresponding TX queue of the network port. This behavior is desirable when tuning for high throughput because the cost of tail pointer updates to both the RX and TX queues can be spread across 32 packets, effectively hiding the relatively slow MMIO cost of writing to the PCIe\* device.

However, this is not very desirable when tuning for low latency, because the first packet that was received must also wait for the other 31 packets to be received. It cannot be transmitted until the other 31 packets have also been processed because the NIC will not know to transmit the packets until the TX tail pointer has been updated, which is not done until all 32 packets have been processed for transmission.

To consistently achieve low latency even under heavy system load, the application developer should avoid processing packets in bunches. The testpmd application can be configured from the command line to use a burst value of 1. This allows a single packet to be processed at a time, providing lower latency, but with the added cost of lower throughput.

## WITHOUT NUMA ENABLED, MY NETWORK THROUGHPUT IS LOW, WHY?

I have a dual Intel® Xeon® E5645 processors 2.40 GHz with four Intel® 82599 10 Gigabit Ethernet NICs. Using eight logical cores on each processor with RSS set to distribute network load from two 10 GbE interfaces to the cores on each processor.

Without NUMA enabled, memory is allocated from both sockets, since memory is interleaved. Therefore, each 64B chunk is interleaved across both memory domains.

The first 64B chunk is mapped to node 0, the second 64B chunk is mapped to node 1, the third to node 0, the fourth to node 1. If you allocated 256B, you would get memory that looks like this:

```
256B buffer
Offset 0x00 - Node 0
Offset 0x40 - Node 1
Offset 0x80 - Node 0
Offset 0xc0 - Node 1
```

Therefore, packet buffers and descriptor rings are allocated from both memory domains, thus incurring QPI bandwidth accessing the other memory and much higher latency. For best performance with NUMA disabled, only one socket should be populated.

## **I AM GETTING ERRORS ABOUT NOT BEING ABLE TO OPEN FILES. WHY?**

As the DPDK operates, it opens a lot of files, which can result in reaching the open files limits, which is set using the `ulimit` command or in the `limits.conf` file. This is especially true when using a large number (>512) of 2 MB huge pages. Please increase the open file limit if your application is not able to open files. This can be done either by issuing a `ulimit` command or editing the `limits.conf` file. Please consult Linux manpages for usage information.

## **VF DRIVER FOR IXGBE DEVICES CANNOT BE INITIALIZED**

Some versions of Linux IXGBE driver do not assign a random MAC address to VF devices at initialization. In this case, this has to be done manually on the VM host, using the following command:

```
ip link set <interface> vf <VF function> mac <MAC address>
```

where <interface> being the interface providing the virtual functions for example, eth0, <VF function> being the virtual function number, for example 0, and <MAC address> being the desired MAC address.

## **IS IT SAFE TO ADD AN ENTRY TO THE HASH TABLE WHILE RUNNING?**

Currently the table implementation is not a thread safe implementation and assumes that locking between threads and processes is handled by the user's application. This is likely to be supported in future releases.

## WHAT IS THE PURPOSE OF SETTING IOMMU=PT?

DPDK uses a 1:1 mapping and does not support IOMMU. IOMMU allows for simpler VM physical address translation. The second role of IOMMU is to allow protection from unwanted memory access by an unsafe device that has DMA privileges. Unfortunately, the protection comes with an extremely high performance cost for high speed NICs.

Setting `iommu=pt` disables IOMMU support for the hypervisor.

**WHEN TRYING TO SEND PACKETS FROM AN APPLICATION TO  
ITSELF, MEANING SMAC==DMAC, USING INTEL(R) 82599 VF  
PACKETS ARE LOST.**

Check on register `LLE (PFVMTXSSW[n])`, which allows an individual pool to send traffic and have it looped back to itself.

## **CAN I SPLIT PACKET RX TO USE DPDK AND HAVE AN APPLICATION'S HIGHER ORDER FUNCTIONS CONTINUE USING LINUX PTHREAD?**

The DPDK's lcore threads are Linux pthreads bound onto specific cores. Configure the DPDK to do work on the same cores and run the application's other work on other cores using the DPDK's "coremask" setting to specify which cores it should launch itself on.



**IS IT POSSIBLE TO EXCHANGE DATA BETWEEN DPDK  
PROCESSES AND REGULAR USERSPACE PROCESSES VIA  
SOME SHARED MEMORY OR IPC MECHANISM?**

Yes - DPDK processes are regular Linux/BSD processes, and can use all OS provided IPC mechanisms.

**CAN THE MULTIPLE QUEUES IN INTEL(R) I350 BE USED WITH  
DPDK?**

I350 has RSS support and 8 queue pairs can be used in RSS mode. It should work with multi-queue DPDK applications using RSS.

## HOW CAN HUGE PAGE-BACKED MEMORY BE SHARED AMONG MULTIPLE PROCESSES?

See the Primary and Secondary examples in the multi-process sample application.

## WHY CAN'T MY APPLICATION RECEIVE PACKETS ON MY SYSTEM WITH UEFI SECURE BOOT ENABLED?

If UEFI secure boot is enabled, the Linux kernel may disallow the use of UIO on the system. Therefore, devices for use by DPDK should be bound to the `vfiopci` kernel module rather than `igb_uio` or `uio_pci_generic`.