



# AMD SP3 Family 17h Models 30h–3Fh Preferred IO Usage Guide

Publication # **56570** Revision: **0.93**

Issue Date: **March 2020**

© 2019–2020 Advanced Micro Devices, Inc. All rights reserved.

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

---

### **Trademarks**

AMD, the AMD Arrow logo, AMD EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Linux is a registered trademark of Linus Torvalds.

## Contents

---

<b>Chapter 1</b>	<b>Introduction.....</b>	<b>5</b>
<b>Chapter 2</b>	<b>Recommended Usage.....</b>	<b>6</b>
2.1	Steps to Optimize I/O Transactions at the PCIe Level .....	6
2.1.1	Consult with Your PCIe Device Vendor .....	6
2.1.2	Enable Preferred IO Mode .....	6



## Revision History

---

Date	Revision	Description
March 2020	0.93	Initial public release.

---

# Chapter 1 Introduction

---

Maintaining coherency of I/O traffic is a requirement to ensure data integrity of all data transfers to and from I/O devices. There are many ways to enforce coherency, and I/O standards have different methods of enforcing it. AMD Infinity Fabric™ technology connects the PCIe subsystem of an AMD EPYC™ processor-based server to the processor cores and memory. The Infinity Fabric technology controls traffic to enforce coherency in the system. With up to 160 PCIe lanes of I/O in an EPYC 7002 series-based two-socket platform, this coherency enforcement can be quite complex.

AMD's Infinity Fabric technology uses more modern methods to maintain coherency. PCIe uses ordering semantics to enforce strict ordering of the data packets to maintain coherency by default. Strict ordering rules prevent a transaction initiated by PCIe endpoint B from completing before a previous transaction initiated by PCIe endpoint A. This can result in endpoint B's transactions being stalled when endpoint A is a lower-performing device. The PCIe specification does, however, have two other options that a PCIe endpoint can enable to allow higher performance of I/O transactions: Relaxed Ordering (RO) and the newer ID-Based Ordering (IDO). Both of these additional methods enable the system to avoid bottlenecks in different manners. AMD has introduced a third one available only with the EPYC 7002 series processors: Preferred IO.

With the increasing adoption of high-bandwidth adapters, allowing traffic to flow out of order and eliminating bottlenecks is becoming critical. Because not all adapters have implemented RO or IDO, AMD created the Preferred IO mode capability within the EPYC 7002 series of processors to better achieve performance expectations. This mode can be enabled for only a single PCIe root complex (RC) and will affect all PCIe endpoints behind that RC. This setting gives priority to the I/O devices attached to the PCIe slot(s) associated with the enabled RC for I/O transactions.

---

## Chapter 2 Recommended Usage

---

AMD has introduced Preferred IO to complement RO and IDO, rather than replace them. As with most performance tuning options, your exact workload may not perform better with Preferred IO enabled but rather with RO or IDO. AMD recommends that you test your workload under all scenarios to find the most efficient setting.

### 2.1 Steps to Optimize I/O Transactions at the PCIe Level

#### 2.1.1 Consult with Your PCIe Device Vendor

Check to see whether your PCIe device can support IDO or RO. IDO is the more optimal method for increasing I/O performance. If your PCIe device does support it, then ensure that you have the appropriate firmware to enable the feature—not all vendors enable these features by default. Consult with the vendor to ensure one or the other is properly enabled.

#### 2.1.2 Enable Preferred IO Mode

If your PCIe device supports IDO or RO, and it is enabled, you can try enabling Preferred IO to determine whether your specific use case realizes an additional performance improvement. If the device does not support either of those features for your high bandwidth adapter, you should take advantage of the Preferred IO mode capability available on the EPYC 7002 series processors.

Steps to take:

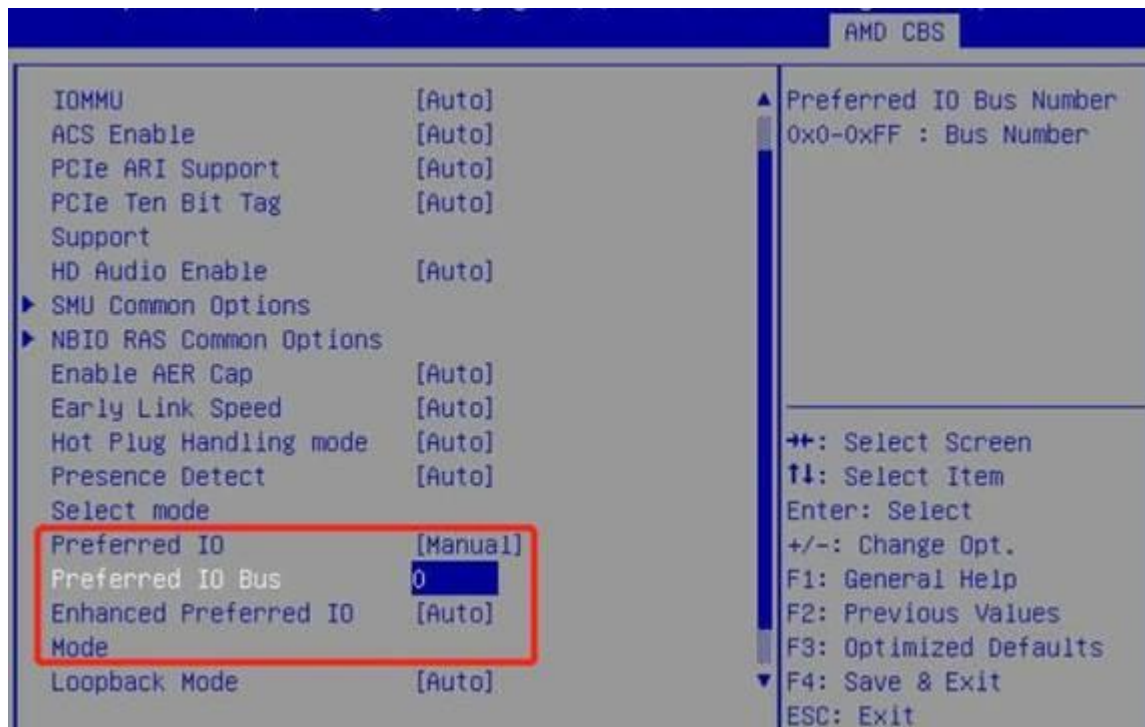
1. Determine at which PCIe bus number your device is located

- For Linux, open a command prompt.
- Type: `lspci` and find the device for which you want to enable Preferred IO mode. One example:

```
81:00.0 SATA controller: Advanced Micro Devices, Inc. [AMD] FCH SATA Controller [AHCI mode] (rev 51)
82:00.0 Non-Essential Instrumentation [1300]: Advanced Micro Devices, Inc. [AMD] Device 1485
82:00.2 Encryption controller: Advanced Micro Devices, Inc. [AMD] Device 1498
83:00.0 Non-Essential Instrumentation [1300]: Advanced Micro Devices, Inc. [AMD] Device 148a
83:00.2 Encryption controller: Advanced Micro Devices, Inc. [AMD] Device 1498
```

- In this example, we want to enable Preferred IO mode on the SATA controller. The output of `lspci` will come in `<bus>:<dev>:<func>` followed by a description provided by the device vendor. We are interested in enabling Preferred IO mode for bus 81.

2. Boot into the system BIOS and enable Preferred IO mode for the slot for your device
  - Each system vendor may locate this option in different locations, but it will typically be found with other I/O specific features.
  - One example location is under *NBIO Common Options* in BIOS:



- After it is enabled, the BIOS setup will ask you for the PCIe *<bus>* which was found with the *lspci* command. For our example, insert *81*. Some OEMs may ask you for a PCIe slot, and they will determine the PCIe bus for you.
  - Note: Earlier versions of BIOS will ask for a device (*<dev>*) and function (*<func>*), but this is not needed by the BIOS code so simply put in 00 for those fields if needed (*81:00:00*).