Falcon Composable Solution

PCIe Resource Evaluation

Version 1.0 October 1st, 2021



© 2021 H3 Platform Inc. or its subsidiaries. All rights reserved.H3 Platform and other trademarks are trademarks of H3 Platform Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Notes, Cautions, and Warning



A NOTE indicates important information that helps you make better use of your product.



A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.



Warning A WARNING indicates a potential for property damage, personal injury, or death.

Note

The information in this document applies only in the context of Falcon composable solutions. The purpose of this document is to help users to evaluate the PCIe resources required on the host machine side prior to implementing Falcon composable solution. Note that the result of evaluation differs according to individual use case.

Host Machine Capability

The number of PCIe devices that a server can handle depends on the PCIe resources that the host can give. The most important parameters in this case are MMIO size and bus numbers, and they vary from server to server. This is because the server vendors design how its BIOS handle the PCIe resources and the different CPU platform possesses different capability. The MMIO size that a server could have depends on the address space of CPU and the limit from BIOS (some BIOS allows maximum 512GB while some allows up to 1TB MMIO), and the bus number allowed usually would not exceed 256 according to PCIe specifications. The host machine should fulfill both MMIO size and bus number requirement when multiple devices are assigned.



Users should make sure that the platform in use is able to provide sufficient PCIe resource to Falcon expansion chassis, including all the peripheral devices installed.

The server will fail to boot when:

- 1. Bus number required exceeds the bus number that the BIOS allows.
- 2. MMIO size required exceed the MMIO size that the CPU and BIOS allow.

Bus Numbers:

An endpoint typically takes up 1 bus, so the more bus numbers that BIOS allows, the more devices can be added to the server theoretically. Note that the bridges also take up buses, and depending on the bridge itself, the number of buses required varies. Note that the bus number required by the bridges are different under standard mode and advanced mode. (Please see user manuals for more info about modes)

(Continue to next page)

1

Table 1. Bus number required (per device drawer) for each Falcon model:

	Bus number required	
Host bus adapter card	13	
Model	Standard Mode	Advanced Mode
Falcon 4005	17	10
Falcon 4010	17	10
Falcon 4205	17	10
Falcon 4210	17	10
Falcon 4109	27	18
Falcon 4118	27	18

E.g., Falcon 4205 standard mode would require a bus number of 30 (13+17) in total.



Bus number requirement - Falcon 4205 standard mode:

For dual-drawer models (4010, 4210, 4118), when connecting both device drawers to the same host machine, the bus number required is doubled.

E.g., Falcon 4210 standard mode with 2 HBA would require a bus number of 60 (2*(13+17)) in total.



MMIO Size:

When booting up a server, the BIOS scans for PCIe devices during booting process, and it will allocate the MMIO resource to all the "bridged" endpoints (peripheral devices that are not directly connected to bus0).

Larger MMIO size usually means that the host could handle more devices. The number of devices depends on the RAM of devices. The key is that MMIO size should be greater than the sum of all device memory.

Minimum MMIO size required can be calculated with following formula:

MMIO size required = Σ endpoint memory size

E.g., if there are 4 GPUs, each with 48 GB RAM, it would require MMIO size of at least 192 GB.



MMIO size required = 48 + 48 + 48 + 48