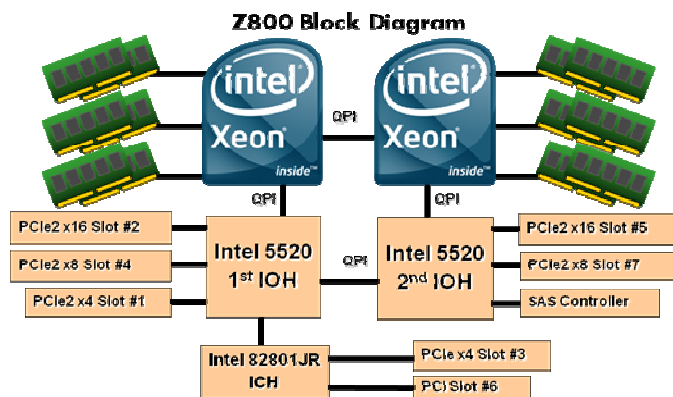


Optimizing PCI-Express Bandwidth on Z800

Introduction: This document provides an overview of how to configure the Z800 to get optimal I/O bandwidth on PCI-Express.

Benefits

For high I/O bandwidth applications, the choice of slot loading, CPU, and memory configuration can be optimized to ensure maximum bandwidth available. Applications and cards sensitive to I/O latency may benefit as well from some of the tips below.



Recommended Configurations Steps

1. Place the graphics cards first. First graphics in slot 2. Second graphics in slot 5.
2. Place I/O cards, from highest bandwidth to lowest, following the slot order listed in the Recommended column of Table 1. This is the optimal load order for most applications.
3. If the onboard SAS controller is not used and there is an I/O card in slot 7, then disable the SAS controller (BIOS setup menu -> Security -> Device Security -> SAS Controller = Device hidden).
4. Additional I/O bandwidth refinements may be possible. If necessary, refer to the tips below.

Table 1. Z800 I/O slot recommended load order

Slot #	PCI-Express		I/O Card Load Order			Bandwidth (GB/s)	IOH
	Lane	Width	Recommended	Opt A	Opt B		
Slot 1	x4	(Gen2)	5 th	5 th	3 rd	2.0	1 st
Slot 2	x16	(Gen2)	1 st	1 st	1 st	8.0	1 st
Slot 3	x4	(Gen1)	6 th	6 th	6 th	1.0	ICH
Slot 4	x8	(Gen2)	3 rd	4 th	2 nd	4.0	1 st
Slot 5	x16	(Gen2)	2 nd	2 nd	4 th	8.0	2 nd
Slot 6	PCI		-	-	-	0.125	ICH
Slot 7	x8	(Gen2)	4 th	3 rd	5 th	4.0	2 nd

Additional Tips for Optimizing PCI-Express I/O Bandwidth

- Select CPUs with the highest QPI frequency (6.4GT/s), especially if dual CPUs are used.
- Make sure all I/O cards are loaded in slots that have a PCI-Express Lane Width at least as wide as the card (see Table 1).
- As with memory performance, NUMA mode may produce inconsistent I/O performance. If this is suspected, try using MNI (Memory Node Interleaved) mode, which has more consistent I/O performance (BIOS setup menu -> Advanced -> Chipset/Memory -> Memory Node Interleave = Enabled).
- Some applications may perform better using one of the optional load orders. For dual CPU systems try load Option A. This may be beneficial for applications which transfer large amounts of data between two I/O cards. For single CPU systems, applications and card sensitive to I/O latency may benefit from load Option B, which will reduce latency by placing cards closer to the CPU.

Links

For information on NUMA, MNI, and Z800 architecture, refer to the corresponding [white papers](#).

