SmartSilc
Solutions by Silicom Ltd.

Boosting Up Intel® QuickAssist Technology 8955 Power with Dual Chip Accelerator Adapter

Mar. 2015

ABSTRACT

The use of hardware accelerator look aside engine often is a good idea for both lower CPU load and freeing up CPU cycles for other business logic processing, and for higher overall performance. It is, however, a path that needs to be taken with some understanding of the server architecture in place, and mainly of the hardware aspects of the data path, in order to gain the best benefit of the acceleration.

1. INTRODUCTION

For network traffic processing with SSL acceleration, where the most common use case is web service termination, then the hardware data paths capabilities are to be accounted for, to calculate correctly the expected results, and to **plan a scale up** of acceleration, when it is possible and required.

The hardware data path can be depicted as a pipeline, or a cascaded funnel, where on each stage in the pipe, on each funnel, has its own capacity of data transfer. Therefore, the path of data transfer from CPU main memory to an acceleration engine (and back, vice versa), is roughly as follows:

1) First funnel – **CPU core** capacity, of handling and processing this much TCP connections per second, or that much data bandwidth transfer;
2) Second funnel – **PCIe interface**, which is the current common interface to look aside engines;
3) Third funnel – **Acceleration chipset** capacity for various types of processing.

The total benefit of the offload and acceleration is a subtotal of those three basic elements.

Two aspects of SSL traffic termination are to be considered when hardware acceleration is evaluated. First, is the number of SSL handshakes per second, where with RSA asymmetric cryptography, the hardware accelerator is used for exponent calculation. Second, is the total bandwidth of data that is transferred to and from acceleration engine, which is the case with symmetric bulk cryptography.

Keywords

SSL Offload, Acceleration, dual chip, PCIe

2. PCIe CONSIDERATIONS

Silicom portfolio of SSL acceleration adapters include, among others, two types of **dual chipset Intel® 8955** based PCIe adapters:

Adapter	PE3iS2CO3X8	PE3iS2CO3
PCIe	Gen 3 x8	Gen 3 x16
RSA 2K / sec	80,000	80,000
AES128	54Gbps	100Gbps

From the above data exemplifies how a hardware accelerator look aside engine should be matched to a specific case and a specific housing.

Let's take the first PCIe x8 based adapter (marked E3iS2CO3X8), for instance. Such an adapter would **fit perfectly for a purpose of accelerating RSA 2K calculations**. Its capacity to uphold up to cool 80 thousand RSA exponents calculations per second, and by that, to support this rate of SSL handshakes, is not hampered by interface bandwidth. Even in this maximal rate, the total bandwidth of data that traverses the PCIe bus to and from the adapters is less than 2Gbps, well within the capacity of x8 lanes of PCIe v3. On the other hand, same bus dictates a hard limit of bandwidth, and thus a hard limit of AES128 symmetric encryption bandwidth.

This limit is immediately overcome with the x16 based adapter (PE3iS2CO3) where PCIe is no longer a bottle neck. It has to be remembered, however, that common commercial off the shelf server often include single PCIe v3 x16 slot and several (six, eight, even more) PCIe v3 x8 slots, a fact that pose physical consideration that is added to the equation.

Silicom ltd., 8 Hanagar St., Kefar Sava 4442537, ISRAEL
sales@silicom.co.il www.silicom-usa.com www.smartsilc.com

2. CPU CONSIDERATIONS

Last, and certainly not least, is the CPU choice. In this context of hardware assisted SSL acceleration, there are two major point to be kept in mind:

1. Number of CPU sockets
2. Number of CPU cores and cores affinity.

While the latter was discussed on another white paper[1], a note on CPU socket would be emphasized herein. It is highly important to avoid the performance degradation that may be incurred to the QPI bus. Although this bus is constantly improving[2], it is best to avoid data transfer from one NUMA[3] node to another.

Therefore, it is advised to connect the hardware accelerator to the same NUMA node (or CPU socket, for that matter), that the network adapter is connected to. That way, offload would serve at its best.

Picture 1 and Picture 2 demonstrate this idea. It may be, in light of the above, that using two dual chipsets adapters, one on each NUMA node, would make more sense than using single quad adapter connected to one NUMA node (through a limited bandwidth PCIe bus); all depends on the specific use case.

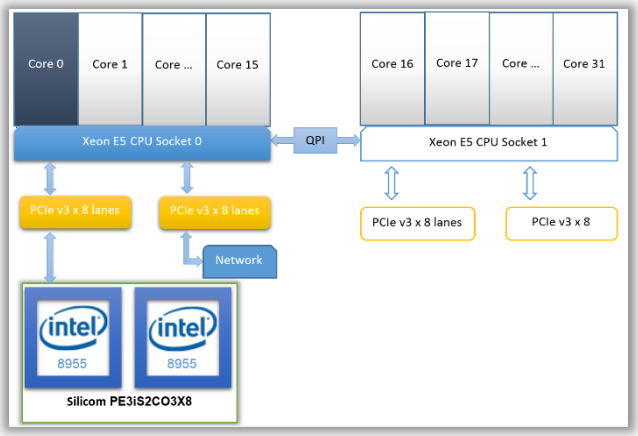

Figure 1 - SIlicom PCIe x8 Based Adapter

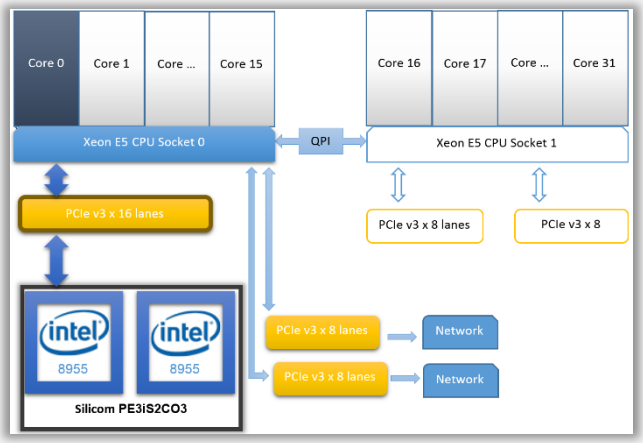

Figure 2- - SIlicom PCIe x16 Based Adapter

REFERENCES

[1] Silicom COWP_FEB2015-v1.1.pdf
[2] 4.0 GHz with Sandy Bridge-E/EP platform and 4.8 GHz with the Haswell-E/EP platform - http://en.wikipedia.org/wiki/Intel_QuickPath_Interconnect
[3] Non-Uniform Memory Access, where memory access time depends on the proximity of the memory to the processor - http://en.wikipedia.org/wiki/Non-uniform_memory_access