



GigalIO™ FabreX™ CPU to CPU Communication

CONTENTS

Introduction.....	3
The Basics of CPU to CPU Communication	3
FabreX Technique for CPU to CPU Communication	4
A Typical Deployment	7
About GigalO Networks.....	8

FIGURES

<i>Figure 1 – Traditional Server Architecture.....</i>	<i>3</i>
<i>Figure 2 – Traditional CPU Architecture.....</i>	<i>4</i>
<i>Figure 3 – Virtual Memory Space</i>	<i>5</i>
<i>Figure 4 – Physical To Virtual Translation</i>	<i>6</i>
<i>Figure 5 – FabreX Deployment.....</i>	<i>7</i>

This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© GigalO Networks, Inc. All rights reserved. GigalO and its affiliates cannot be responsible for errors or omissions in typography or photography. GigalO, the GigalO logo, and FabreX are trademarks of GigalO Networks, Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. GigalO disclaims proprietary interest in the marks and names of others.

June 2019 | Rev 1.0



Introduction

This Application Note is intended to describe CPU to CPU communication with FabreX.

This discussion pertains to any resource that has some sort a processing element in it with a supporting (RCP) Root Complex Processor that is able to enumerate PCIe (PCI Express) I/O devices. Consequently, resources such as FPGA, AI ASICS, intelligent storage, embedded controller, etc. can all potentially fall under this category labeled 'CPU'.

The Basics of CPU to CPU Communication

The discussion in this document is confined to traditional CPU architecture as prevalent in servers. However, it can be extrapolated to other intelligent agents alluded to in the previous section.

Figure 1 depicts such a traditional architecture, where it shows 'n' number of servers communicating amongst themselves. The dotted lines show the communication paths from a conceptual standpoint.

This figure also shows the fundamental architecture of the CPU ecosystem present in practically all servers.

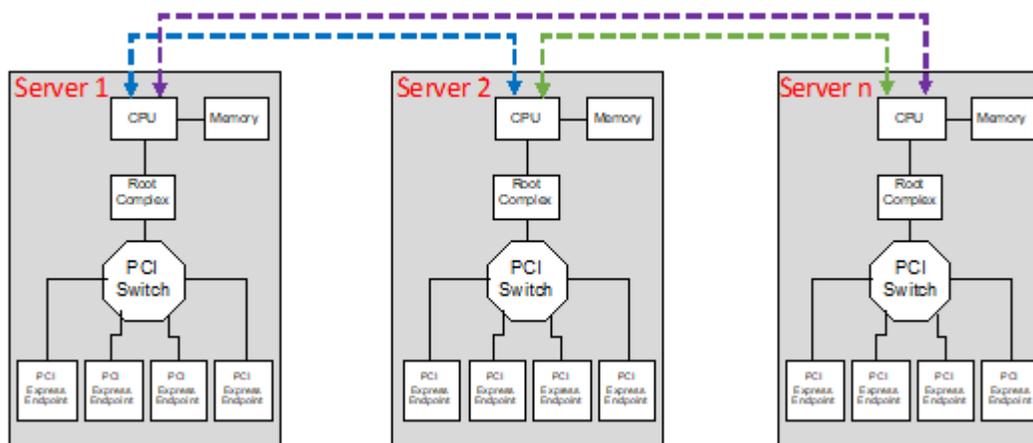


FIGURE 1 – TRADITIONAL SERVER ARCHITECTURE

CPUs do not have the capability to directly read/write each other resulting in the inability to directly communicate with each other.

The memory associated with the individual CPU is what is used as an intermediary resource to pass messages between CPU. There are also sophisticated techniques to inform CPUs of pending messages waiting to be serviced that originated from another CPU. In PCIe technology this mechanism is a vectored interrupt operation consisting of a memory write to a specific memory address location to the receiving agent by the sending agent.



Figure 2 shows the traditional CPU architecture in server environments. CPUs generally have multiple cores with their own exclusive L1 and L2 cache that simultaneously execute individual threads to boost the overall throughput of the CPU.

A point to note here is that there is a PCIe interface to the system memory consisting of DDR3/4 memory devices as well as the L3 cache that can be shared by all the cores. This also shows the entire System memory of the CPU as well as the L3 cache mapped within the PCIe memory space of every server via the built-in PCIe interface.

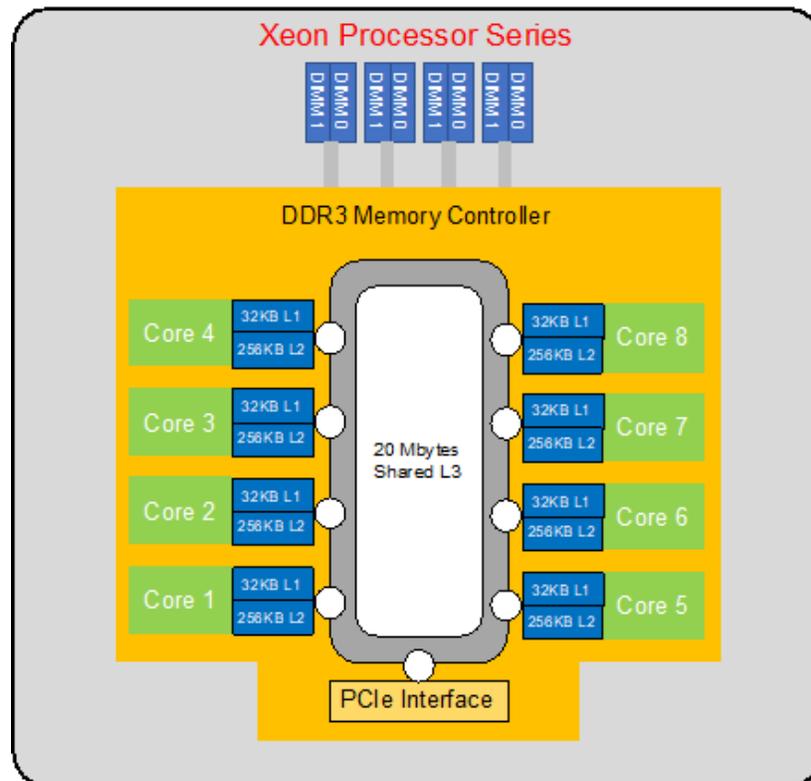


FIGURE 2 – TRADITIONAL CPU ARCHITECTURE

FabreX Technique for CPU to CPU Communication

FabreX uses a technique whereby the FabreX switch controller under the direction of and in cooperation with the Operating System of the server maps a segment of the pertinent system memory physical space reserved for communication to a virtual address space within the FabreX domain.

Figure 3 shows an example of this. All of the individual system memory windows of the respective servers as depicted in **Figure 1** are assigned a memory window of the same size in the 64-bit Global Virtual Address Space of FabreX, i.e. there is no physical memory resident in this space. The memory window sizes shown of the respective servers can be different from each other and mapped to any arbitrary physical address of the server's system memory. These



parameters are coordinated between the OS running on the various Servers and the FabreX Switch Controller resident in the Switch.

One point to note is that these windows in the Virtual Address space are absolutely unique with no overlap whatsoever.

These Virtual Memory windows do not necessarily have to correspond to server system physical memory but instead can be a block of physical memory that can be shared by all servers attached to FabreX.

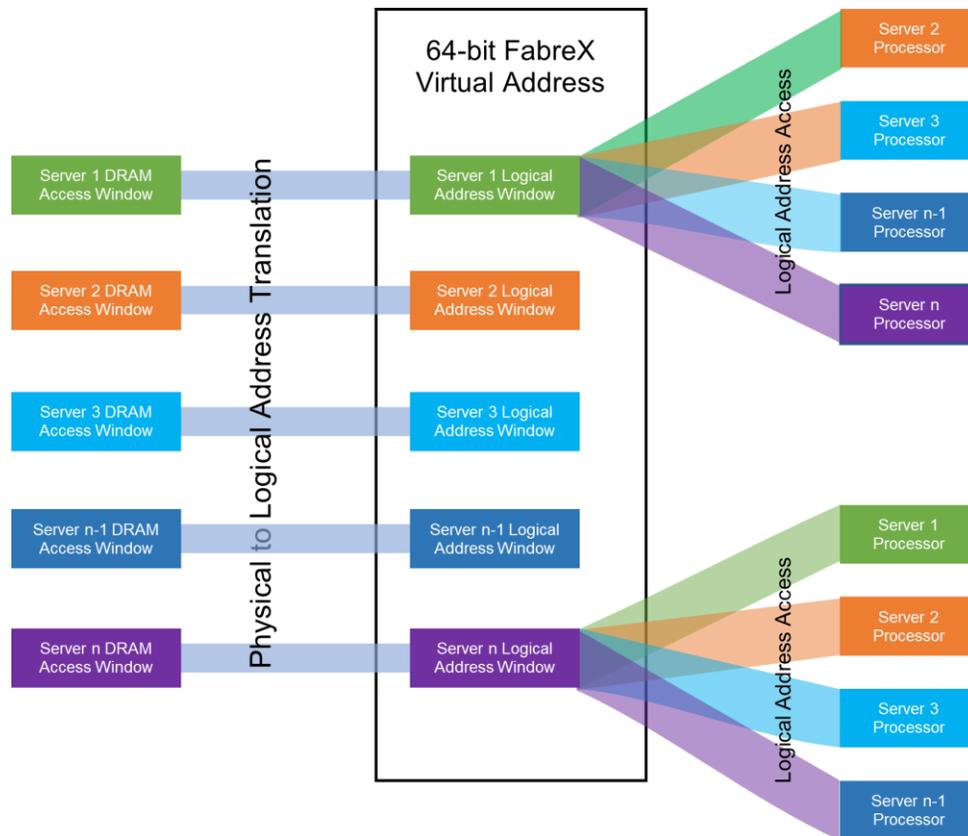


FIGURE 3 – VIRTUAL MEMORY SPACE

Figure 4 elaborates on this concept. This shows a non-volatile memory constituting 3D-XPoint technology as the shared memory for this architecture. This was selected on account of this being cell addressable non-volatile memory with fast access times, and unlike NVMe, it is not block addressable.

This diagram depicts how a FabreX centric system is architected for CPU to CPU communication.

This shows every server has a space dedicated in its physical memory that is accessible by all other servers attached to FabreX. Obviously there needs to be memory protection and some



rudimentary protocols to prevent over writing of the message or communication space, whose implementation falls under the domain of the specific applications running on the servers. This figure highlights one approach of communicating between servers where the servers system address space is divided into physical memory windows supported by memory hardware devices and logical memory space where the physical memory devices are located remotely. All of these spaces fall under the PCIe addressable space of every individual server. Again, all of these assignments are coordinated and implemented by the Servers and the Switch Controller at Boot time.

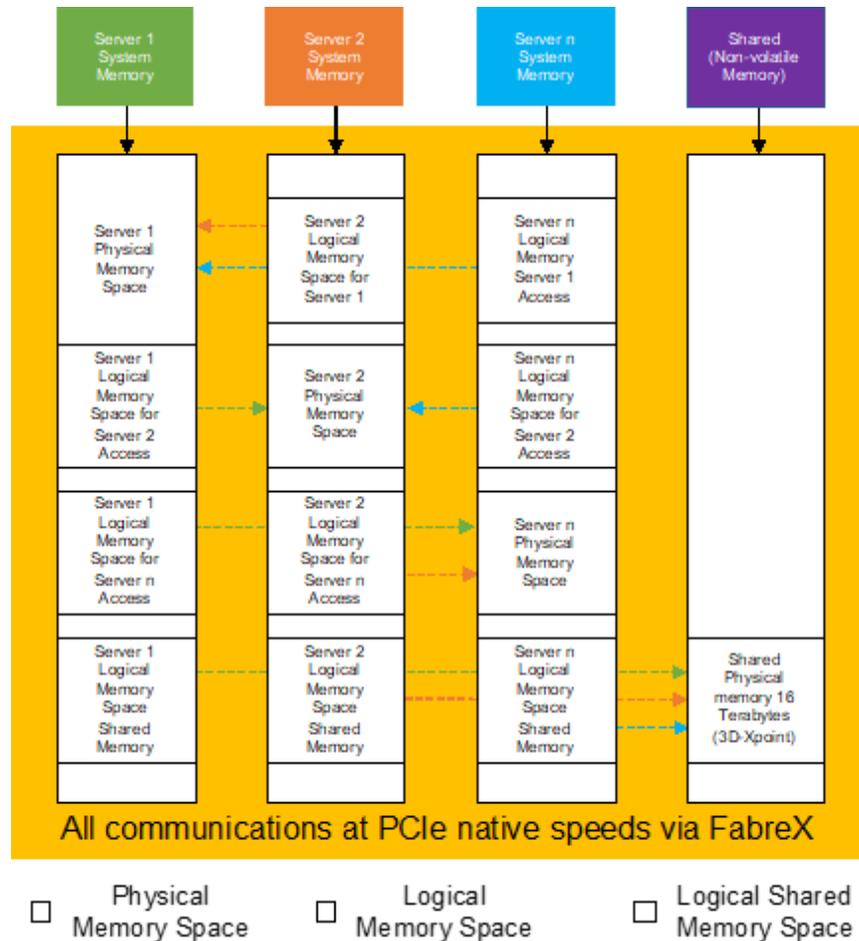


FIGURE 4 – PHYSICAL TO VIRTUAL TRANSLATION

Incidentally, in some implementation the entire system memory of all servers can be made to be visible to all other servers. This would be an ideal solution as long as the software running in individual servers were disciplined to not access protected memory regions of system memory of the individual servers.



A Typical Deployment

Figure 5 depicts a deployment of this architecture with FabreX.

It should now be obvious from the discussion above that FabreX is truly a memory-centric fabric. Communications between agents boil down to the very fundamental rudimentary operations of ‘Load’ and ‘Store’. In addition, this communication path is bolstered by using the robust PCIe protocol, soon to be CXL, both delivering the lowest possible latency in the industry.

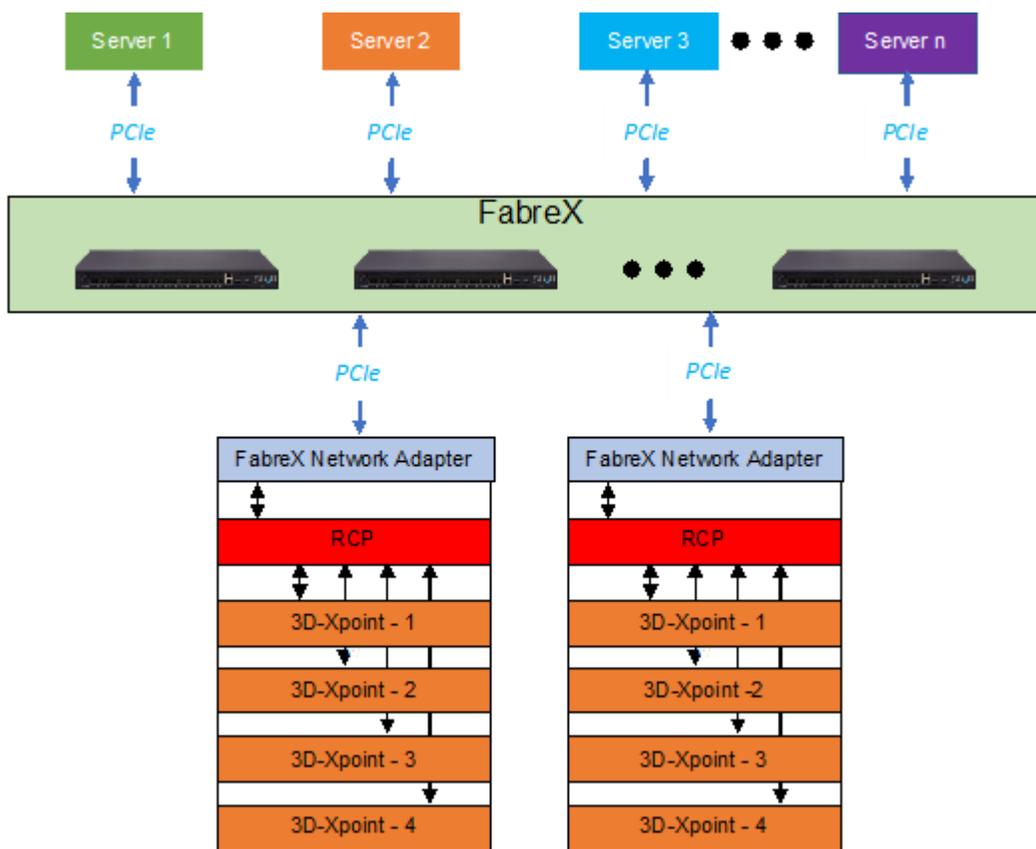


FIGURE 5 – FABREX DEPLOYMENT



About GigaIO Networks

Headquartered in Carlsbad, California, GigaIO democratizes AI and HPC architectures by delivering the elasticity of the cloud at a fraction of the TCO (Total Cost of Ownership). With its universal dynamic infrastructure fabric, FabreX™, and its innovative open architecture using industry-standard PCI Express/soon CXL technology, GigaIO breaks the constraints of the server box, liberating resources to shorten time to results. Data centers can scale up or scale out the performance of their systems, enabling their existing investment to flex as workloads and business change over time. For more information, contact info@gigaio.com or visit www.gigaio.com. Follow GigaIO on [Twitter](#) and [LinkedIn](#).

