

Entering a New Arena in Computing

By: Munira Hussain, December 2016 – HPC Innovation Lab

In order to build a balanced cluster ecosystem and eliminate bottle-necks, the need for powerful and dense server node configurations is essential to support parallel computing. The challenge is to provide maximum compute power with efficient I/O subsystem performance, including memory and networking. Some of the emerging technologies along with traditional computing that are needed for intense compute power are advanced parallel algorithms in the areas of research, life science and financial application.

Dell PowerEdge C6320p

The introduction of the Dell EMC C6320p platform, which is one of the densest and greatest maximum core capacity platform offerings in HPC solutions, provides a leap in this direction.

The PowerEdge C6320p platform is Dell EMC's first self-bootable Intel Xeon Phi platform. The previously available versions of Intel Xeon Phi were PCIE adapters that required a host system to be plugged into. From the core perspective, it supports up to 72 processing cores, with each core supporting two vector processing units capable of AVX-512 instructions. This increases the computation of floating point operations requiring longer vector instructions unlike Intel Xeon® v4 processors that support up to AVX-2 instructions. The Intel Xeon Phi in Dell EMC C6320p also features on-package 16GB of fast MCDRAM that is stacked on the processor. The availability of MCDRAM helps out-of-order execution in applications that are sensitive to high memory bandwidth. This is in addition to the six channels of DDR4 memory hosted on the server. Being a single socket server, the C6320p provides a low power consumption compute node compared to traditional two socket nodes in HPC.



The following table shows platform differences as we compare the current Dell EMC PowerEdge C6320 and Dell EMC PowerEdge C6320p server offerings in HPC.

	C6320	C6320p
Server Form Factor	2U Chassis with four sleds	2U Chassis with four sleds
Processors		
	Dual Socket	Single Socket
Processor Model	Intel® Xeon	Intel® Xeon Phi
Max cores in a sled	Up to 44 physical cores, 88 logical cores (with two * Intel® Xeon E5-2699 v4, 2.2 GHz, 55MB, 22 cores, 145W)	Up to 72 physical cores, 288 logical cores (with the Intel® Xeon Phi Processor 7290 (16GB, 1.5GHz, 72 core, 245W))
Theoretical DP Flops per sled	1.26 TFlops/sec	2.9 TFlops/sec
Memory		
DIMM slots	16 DDR4 DIMM slots	6 DDR4 DIMM slots + on-die 16GB MCDRAM
MCDRAM BW (Memory mode)	N/A	~ 475-490 GB/s
DDR4 BW	~ 135 GB/s	~ 90GBs
Interconnects		
Provisioning fabric	Dual port 1Gb/10GbE	Single port 1GbE
High Speed	Intel Omni-Path Fabric (100Gbps) Mellanox Infiniband (100Gbps)	Intel Omni-PathFabric (100Gbps) On-board Mellanox Infiniband (100Gbps)
Storage		
	Up to 24 x 2.5" or 12 x 3.5" HD	6 x 2.5" HD per node + Internal 1.8" SSD option for boot
Management		

Integrated Dell EMC Remote Access Controller	Dedicated and shared iDRAC8	Dedicated and shared iDRAC8
---	-----------------------------	-----------------------------

Table 1: Comparing the C6320 and C6320p offering in HPC

Dell EMC Supported HPC Solution:

Dell EMC offers a complete, tested, verified and validated solution offering on the C6320p servers. This is based on Bright Cluster Manger 7.3 with RHEL 7.2 that includes specific highly recommended kernel and security updates. It will also provide support for the upcoming RHEL 7.3 operating system. The solution provides automated deployment, configuration, management and monitoring of the cluster. It also integrates recommended Intel performance tweaks, as well as required software drivers and other development toolkits to support the Intel Xeon Phi programming model.

The solution provides the latest networking support for both InfiniBand and Intel Omni-Path Fabric. It also includes Dell EMC-supported System Management tools that are bundled to provide customers with the ease of cluster management on Dell EMC hardware.

**Note: As a continuation to this blog, there will be follow-on micro-level benchmarking and application study published on C6320p.*

OS name: Linux OS version: 3.10.0-327.36.3.el7.x86_64 OS flavor: RedHat7 u2 Motherboard manufacturer: Dell Inc. Motherboard name: 0D2RR8 Bios vendor: Dell Inc. System manufacturer: Dell Inc. System name: PowerEdge C6320p Vendor tag: Not Specified Bios date: 12/001/2016 Bios version: 1.1.3		Disk Information <table border="1"> <thead> <tr> <th>Name</th> <th>Model</th> <th>Size</th> </tr> </thead> <tbody> <tr> <td>sdb</td> <td>ST91000640NS</td> <td>931.51 GiB</td> </tr> <tr> <td>sda</td> <td>MZ8LM120HCFD00D3</td> <td>111.79 GiB</td> </tr> </tbody> </table>	Name	Model	Size	sdb	ST91000640NS	931.51 GiB	sda	MZ8LM120HCFD00D3	111.79 GiB																																														
Name	Model	Size																																																							
sdb	ST91000640NS	931.51 GiB																																																							
sda	MZ8LM120HCFD00D3	111.79 GiB																																																							
CPU Information <table border="1"> <thead> <tr> <th>CPU id</th> <th>Core id</th> <th>Model name</th> <th>Speed</th> <th>Cache</th> <th>Current governor (others)</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1498 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> <tr> <td>0</td> <td>1</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1277 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> <tr> <td>0</td> <td>2</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1498 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> <tr> <td>0</td> <td>3</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1504 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> <tr> <td>0</td> <td>4</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1500 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> <tr> <td>0</td> <td>5</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1505 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> <tr> <td>0</td> <td>6</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1302 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> <tr> <td>0</td> <td>7</td> <td>Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz</td> <td>1506 MHz</td> <td>1024 KB</td> <td>performance (powersave)</td> </tr> </tbody> </table>				CPU id	Core id	Model name	Speed	Cache	Current governor (others)	0	0	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1498 MHz	1024 KB	performance (powersave)	0	1	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1277 MHz	1024 KB	performance (powersave)	0	2	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1498 MHz	1024 KB	performance (powersave)	0	3	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1504 MHz	1024 KB	performance (powersave)	0	4	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1500 MHz	1024 KB	performance (powersave)	0	5	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1505 MHz	1024 KB	performance (powersave)	0	6	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1302 MHz	1024 KB	performance (powersave)	0	7	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1506 MHz	1024 KB	performance (powersave)
CPU id	Core id	Model name	Speed	Cache	Current governor (others)																																																				
0	0	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1498 MHz	1024 KB	performance (powersave)																																																				
0	1	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1277 MHz	1024 KB	performance (powersave)																																																				
0	2	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1498 MHz	1024 KB	performance (powersave)																																																				
0	3	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1504 MHz	1024 KB	performance (powersave)																																																				
0	4	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1500 MHz	1024 KB	performance (powersave)																																																				
0	5	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1505 MHz	1024 KB	performance (powersave)																																																				
0	6	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1302 MHz	1024 KB	performance (powersave)																																																				
0	7	Intel(R) Xeon Phi(TM) CPU 7290F @ 1.50GHz	1506 MHz	1024 KB	performance (powersave)																																																				

References:

- 1) http://www.dell.com/us/business/p/poweredge-c6320p/pd?ref=PD_OC
- 2) http://en.community.dell.com/techcenter/high-performance-computing/b/general_hpc/archive/2016/11/11/cryo-em-in-hpc-with-knl