Dell EMC HPC Systems - SKY is the limit

Munira Hussain, HPC Innovation Lab, July 2017

This is an announcement about the Dell EMC HPC refresh that introduces support for 14th Generation servers based on the new Intel® Xeon® Processor Scalable Family (micro-architecture also known as “Skylake”). This includes the addition of PowerEdge R740, R740xd, R640, R940 and C6420 servers to the portfolio. The portfolio consists of fully tested, validated, and integrated solution offerings. These provide high speed interconnects, storage, an option for both hardware and cluster level system management, and monitoring software.

On a high level, the new generation Dell EMC Skylake servers for HPC provide greater computation power, which includes support for up to 28 cores and memory speed up to 2667 MT/s; the architecture extends AVX instructions to AVX512. The AVX512 instructions can execute up to 32 DP FLOP per cycle, which is twice the capability of the previous 13th generation servers that used Intel Xeon E5-2600 v4 processors (“Broadwell”). Additionally, the number of core counts per socket is 20% higher per system when compared to the previous generation, which consisted of a maximum 22 cores. It consists of six memory channels per socket; therefore, a minimum of 12 DIMMs are needed for a dual socket server to provide up to full memory bandwidth. The chipset also has 48 PCI-E lanes per socket, up from 40 lanes in the previous generation.

The table below notes the enhancements in the latest PowerEdge servers over the previous generations:

High Level Comparison of the Dell EMC Server Generations for HPC Offering:

	PowerEdge R730 (13th Generation Server)	PowerEdge R740 (14th Generation Server)
Chipset		

CPU and Chipset	Grantly with Broadwell –EP	Purley with Skylake
PCH	Wellsburg DMI2	Lewisburg DMI3
Front Side Bus	Intel® QuickPath Interconnect	Intel® UltraPath Interconnect
Processor and cores		
Cores (max)	22 cores	28 cores
TDP (max)	145W	205W
Instruction Set	AVX2	AVX2/AVX-512
Max DP FLOPS/CLK	16 per core (w/AVX2)	32 per core (w/AVX-512)
Memory		
Memory Channels per Socket	4 channels, DDR4	6 channels, DDR4
Memory Speed (upto max speed)	1DPC=2400, 2DPC =upto 2400, 3DPC =1866	1 DPC and 2DPC =2133, 2400, 2666 No 3DPC
NVDIMM	N/A	Up to 12 NVDIMM
GPGPU		
GPGPU x16 PCIE Risers	Up to Two Double width	Up to Three Double width
Storage		
PERC	PERC 9	PERC 9/10
HardDrive	SATA/SAS/SSD	SATA/SAS/SSD
		Boot Optimized Storage Subsystem (BOSS)
		Support for NVMe
Dell EMC Hardware System Management		
System Management	iDRAC8	iDRAC9
IDRAC IP Source (Factory default)	Static IP	DHCP
IDRAC Login (Factory default)	Standard	Random for enhanced security
System Management tools	OpenManage 8.5 DTK 5.5	OpenManage 9.0.1 DTK 6.0.1
BIOS		
Default mode	Legacy	UEFI
Performance		

Max Total System Memory BW (STREAM TRIAD)	130 (GBps) 1DCP with 2400 MT/s	215 (Gbps) 2DPC with 2666 MT/s
Theoretical max System Performance (HPL) ***This is based on the max core – memory configuration***	1.4 TFlops (Dual Socket – 22 c , 145W, 2400 MT/s) based on AVX2	3.2 Tflops (Dual Socket – 28 c, 205W, 2666 MT/s) Based on AVX-512 code

The HPC release supporting Dell EMC 14G servers is based on the Red Hat Enterprise Linux 7.3 operating system. It is based on the 3.10.0-514.el7.x86_64 kernel. The release also supports the new version of Bright Cluster Manager 8.0. Bright Cluster Manager (BCM) is integrated with Dell EMC-supported tools, drivers, and third-party software components for the ease of deployment, configuration, and management of the cluster. It includes Dell EMC System Management tools based on OpenManage 9.0.1 and Dell EMC Deployment ToolKit 6.0.1 that help manage, monitor, and administer Dell EMC hardware. Additionally, updated third party drivers and development tools from Mellanox OFED for InfiniBand, Intel IFS for Omni-Path, NVIDIA CUDA for latest Accelerators, and other packages for Machine Learning are also included. Details of the components are as below:

- Based on Red Hat Enterprise Linux 7.3 (kernel 3.10.0-514.el7.x86_64)
- Dell EMC System Management tools from Open Manage 9.0.1 and DTK 6.0.1 for 14G and Open Manage 8.5 and DTK 5.5 for up to 13G Dell EMC servers
- Updated Dell EMC supported drivers for network and storage deployed during install
  - megaraid_sas = 7.700.50
  - igb=5.3.5.7
  - ixgbe=4.6.3
  - i40e=1.6.44
  - tg3=1.137q
  - bnx2=2.2.5r
  - bnx2x=1.714.2
- Mellanox OFED 3.4 and 4.0 for InfiniBand
- Intel IFS 10.3.1 drivers for Omni-Path
- CUDA 8.0 drivers for NVidia accelerators
- Intel XPPSL 1.5.1 for Intel Xeon Phi processors

- Additional Machine Learning packages such as TensorFlow, Caffe, Cudnn, Digits and required dependencies are also supported and available for download

Below are some images of the Bright Cluster Manager 8.0 BrightView:

Figure1: This shows the overview of the Cluster. It displays the total capacity, usage, and job status.

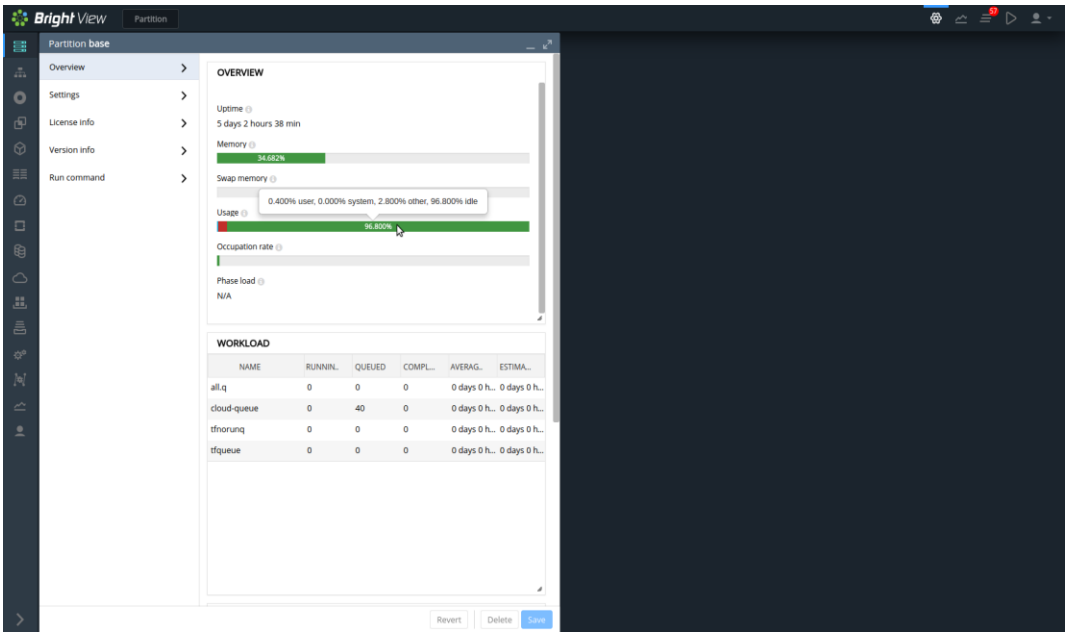

Figure 2: Displays the cascading view of Cluster configuration and respective settings within a group. The settings can be modified and applied from the console.

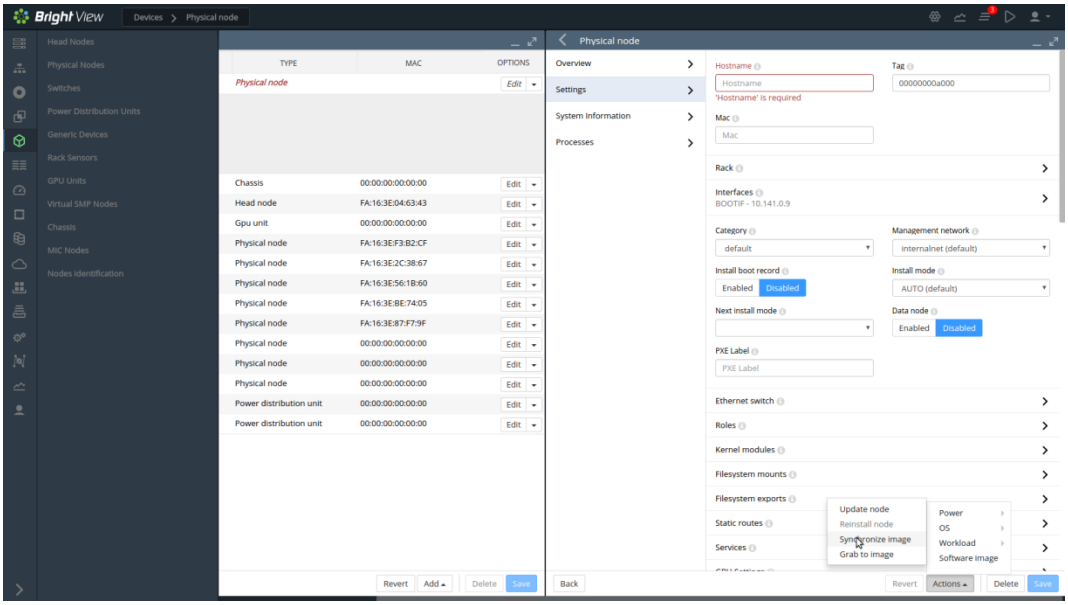

Figure 3: Dell EMC Settings Tab shows the parsed info on hardware configuration and the required BIOS level settings.

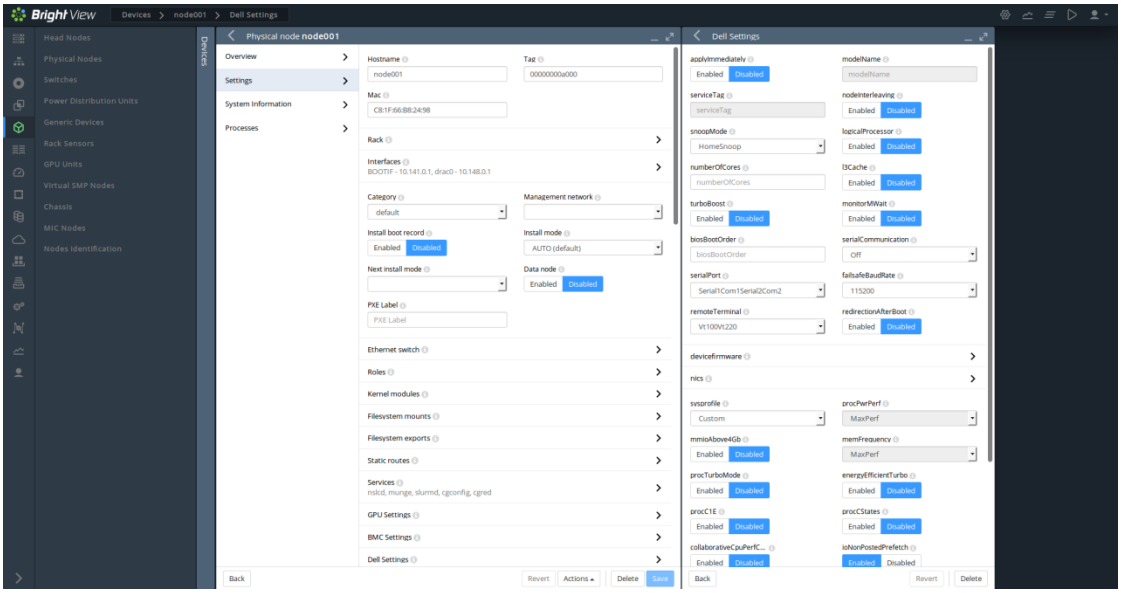

Dell EMC HPC Systems based on the 14th Generation servers expand HPC computation capacity and demands. They are fully balanced and architected solutions that are validated and verified for the customers, and the configurations are scalable. Please stay tuned as follow-on blogs will cover performance and application study; these will be posted here: http://en.community.dell.com/techcenter/high-performance-computing/