

Reference Architecture of Dell EMC Ready Solution for HPC Life Sciences

Refresh with Cascade Lake CPUs and Dell EMC Ready Solution for HPC Storage

Abstract

The Dell EMC's flexible HPC architecture for Life Sciences has been through a dramatic improvement with new Intel® Xeon® Scalable Processors. Dell EMC Ready Bundle for HPC Life Sciences equipped with better 14G servers, faster CPUs, and more memory bring a much higher performance in terms of throughput compared to the previous generation especially in genomic data processing.

October 2019

Revisions

Date	Description
March 2018	Initial release with Intel® Xeon® Scalable processors (code name Skylake)
October 2019	Revised with 2nd Generation Intel Xeon Scalable processors (code name Cascade Lake)

Acknowledgements

This paper was produced by the following:

Author: Kihoon Yoon

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [11/14/2019] [Reference Architecture] [Document 309]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	4
1 Introduction.....	5
2 Solution Overview	6
2.1 Architecture.....	6
2.2 Compute and Management Components	7
2.3 Storage Components.....	9
2.3.1 Dell EMC Ready Solution for HPC Lustre Storage	9
2.4 Network Components	11
2.4.1 Dell Networking H1048-OPF	11
2.4.2 Mellanox SB7700.....	11
2.4.3 Dell Networking S3048-ON	11
2.5 Reference Architecture	11
2.6 Software Components	13
3 Performance Evaluation and Analysis	15
3.1 Variant Calling analysis performance	15
3.1.1 Single Sample Test	16
3.1.2 Single Sample per node/Multiple node test.....	17
3.2 <i>De Novo</i> Assembly Performance	18
3.2.1 SOAPdenovo2.....	19
3.2.2 SPAdes.....	19
4 Conclusion.....	21
A Technical support and resources	22
A.1 Related resources.....	22

Executive summary

Since Dell EMC announced the Dell EMC Ready Solution for HPC Life Science in March 2018, the solution has matured, and can now process roughly 150 50x whole human genomes per day using 64x C6420s and Dell EMC Ready Solution for HPC Lustre Storage. This is roughly a two-fold improvement in throughput from the previous solution, made possible by the introduction of Dell EMC's 14th Generation servers which include the latest Intel Xeon Scalable processors, updated server portfolio, improved memory, and storage subsystem performance (1).

This reference architecture paper describes the architectural changes and updates to the follow-on of Dell EMC Ready Solutions for HPC Life Science. It explains new features, demonstrates the benefits, and shows the improved performance of the current offering.

AUDIENCE

This document is intended for organizations interested in accelerating genomic research with advanced computing and data management solutions. System administrators, solution architects, and others within those organizations constitute the target audience.

1 Introduction

Dell EMC has updated PowerEdge 14th Generation servers with 2nd Gen Intel Xeon Scalable processors, a boon to HPC applications. In addition to offering more cores, the solution offers Optane support, faster DRAM (DDR4-2933 in 1 DPC configuration), and more DRAM configurations (1TB, 2TB, and 4TB). Consumers are generally expecting more performance, better efficiency, and lower power from a newer processor. However, some customers look for the improvements which are not obvious such as support for new instructions, layered ecosystem optimizations, support for new technology, or a new product direction. Since 2nd generation processor was built on a foundation of the predecessor focusing on the secondary characteristics, most open-source applications in Life Sciences take a longer time to adopt those advantages. Nonetheless, this reference architecture illustrates how the new CPUs behave on two different genomics workloads, variant calling and *de novo* assembly, especially with the newer version of Dell EMC Ready Solution for HPC Lustre Storage.

2 Solution overview

HPC in Life Sciences requires a flexible architecture to accommodate various system requirements. The Dell EMC Ready Solution for HPC Life Sciences was created to meet this need. It is a pre-integrated, tested, tuned, and leverages the most relevant of Dell EMC's high-performance computing line of products and best-in-class partner products (2). The solution encompasses all the hardware resources required for various life sciences data analysis, while providing an optimal balance of compute density, energy efficiency, and performance.

2.1 Architecture

The Dell EMC Ready Solution for HPC Life Sciences provides high flexibility to cover various needs from the large number of applications. The platform is available in three variants which are determined by the cluster interconnect option selected as the storage component. This reference architecture uses a version of the Lustre/ME4 series to discuss the updated solution. In the current version, the following options are available:

- Dell EMC PowerEdge C6420 compute subsystem with Intel Omni-Path (OPA) fabric or Mellanox InfiniBand (IB) EDR fabric (3)
 - Storage option:
 - > Dell EMC Ready Solution for HPC Lustre Storage as a performance scratch space
 - > Dell EMC Ready Solution for HPC NFS Storage (NSS7.0-HA) as a home directory space¹
- PowerEdge C6420 compute subsystem with 10/40GbE fabric
 - Storage option:
 - > Dell EMC Isilon F800 as a performance scratch space² (4)
 - > Dell EMC Ready Solution for HPC NFS Storage (NSS7.0-HA) as a home directory space
- Add-on compute nodes:
 - Dell EMC PowerEdge R940 (5)
 - > This server covers large memory applications such as *de novo* assembly
 - Dell EMC PowerEdge C4140 (6)
 - > A server for accelerators like NVIDIA GPUs for molecular dynamics simulations
 - Dell EMC PowerEdge R740xd (7)
 - > FPGA and other accelerators can be accommodated with this server.

¹ The current version of Dell EMC Ready Solution for HPC NFS Storage is NSS7.4-HA. NSS7.4-HA will be added during the next iteration.

² There is no change in Dell EMC Isilon F800 for this reference architecture update. Detailed information about this storage option can be found in the previous reference architecture paper (https://downloads.dell.com/manuals/all-products/esuprt_software/esuprt_it_ops_datcentr_mgmt/high-computing-solution-resources_white-papers85_en-us.pdf).

In addition to compute, network, and storage options, there are several other components that perform different functions in the Dell EMC Ready Solution for HPC Life Sciences. These include CIFS gateway, fat node, acceleration node and other management components. Each of these components is described in detail in the subsequent section.

The solutions are nearly identical for Intel OPA and IB EDR versions except for a few changes in the switching infrastructure and network adapter. The solution ships in a deep and wide 48U rack enclosure, which helps to make PDU mounting and cable management easier. In section 2.5, the components of two fully loaded racks using 56x Dell EMC PowerEdge C6420 rack servers as a compute subsystem, a Dell EMC PowerEdge R940 as a fat node, a Dell EMC PowerEdge C4140 as an accelerator node, a Dell EMC Ready Solution for HPC NFS Storage, a Dell EMC Ready Solution for HPC Lustre Storage and an Intel OPA as the cluster's high-speed interconnect.

2.2 Compute and management components

There are several considerations when selecting the servers for master node, login node, compute node, fat node and accelerator node. For master node, 1U form factor Dell EMC PowerEdge R440 is recommended. The master node is responsible for managing the compute nodes and optimizing the overall compute capacity. The login node (Dell EMC PowerEdge R640 is recommended) is used for user access, compilations and job submissions. Usually, master and login nodes are the only nodes that communicate with the outside world, and they act as a middle point between the actual cluster and the outside network. For this reason, high availability (HA) can be provided for master and login nodes as an option.

Ideally, the compute nodes in a cluster should be as identical as possible since the performance of parallel computation is bounded by the slowest component in the cluster. Heterogeneous clusters do work, but careful execution is required to achieve the best performance. For Life Sciences applications, however heterogeneous clusters work well to handle completely independent workloads such as DNA-Seq, *de novo* assembly or molecular dynamics simulations. Because these workloads require quite different hardware components, we recommend Dell EMC PowerEdge C6420 as a compute node to handle NGS data processing due to its density, a wide choice of CPUs, and high maximum memory capacity. Dell EMC PowerEdge R940 is an optional node with 6TB of RDIMM/LRDIMM memory and is recommended for customers who need to run applications requiring large memory such as *de novo* assembly. Accelerators are used to speed up computationally intensive applications such as molecular dynamics simulation applications. We tested configurations K and M for this solution.

The compute and management infrastructure consist of the following components.

- Compute
 - Dell EMC PowerEdge C6400 enclosure with 4x C6420 servers
 - > CPU: 2x Intel Xeon Gold 6248 at 2.4 GHz 20 cores
 - > Memory: 12 x 32 GB at 2933 MHz
 - > Disk: 6x 480 GB 12 Gbps SAS mixed use SSDs
 - > Interconnect: Intel® Omni-Path, IB EDR or 10/40 GbE, or both
 - > BIOS system profile: Performance optimized
 - > Logical processor: Disabled
 - > Virtualization technology: Disabled
 - > Node interleaving: Enabled
 - > Operating system: RHEL 7.6

- Dell EMC PowerEdge C4140 with up to four NVIDIA Tesla V100 GPUs
 - > CPU: 2x Intel Xeon Gold 6148 at 2.4 GHz 20 cores
 - > Memory: 24x 16 GB at 2666 MHz
 - > GPU: 4x V100-SXM2 16 GB
 - > Disk: 9 TB HDD
 - > BIOS system profile: Performance optimized
 - > Logical processor: Disabled
 - > Virtualization technology: Disabled
 - > Operating system: RHEL 7.6
- Dell EMC PowerEdge R940
 - > CPU: 4x Intel Xeon Platinum 8280M at 2.7 GHz 28 cores
 - > Memory: 48x 32 GB at 2666 MHz
 - > Disk: 12x 1.92 TB SSD SATA mixed use 6 Gbps 512e 2.5 in hot plug S4610 drive in RAID 0
 - > BIOS system profile: Performance optimized
 - > Logical processor: Disabled
 - > Virtualization technology: Disabled
 - > Operating system: RHEL 7.6
- Management
 - Dell EMC PowerEdge R440 for master node
 - > CPU: 2x Intel Xeon Gold 6230 at 2.1 GHz 20 cores (Cascade Lake)
 - > Memory: 12x 16 GB at 2666 MHz
 - > Disk: 10x 1.92 TB SSD SATA read intensive 6 Gbps 512e 2.5 in hot-plug S4510 drive in RAID 10
 - > BIOS system profile: performance optimized
 - > Logical processor: Enabled
 - > Virtualization technology: disabled
 - > Operating system: RHEL 7.6
 - Dell EMC PowerEdge R640 for login node and CIFS gateway (optional)
 - > CPU: 2x Intel Xeon Gold 6248 at 2.5 GHz 20 cores
 - > Memory: 12x 16 GB at 2933 MHz
 - > Disk: 10x 1.92 TB SSD SATA mixed use 6 Gbps 512e 2.5in hot plug S4610 drive in RAID 10
 - > BIOS system profile: Performance optimized
 - > Logical processor: Enabled
 - > Virtualization technology: Disabled
 - > Operating system: RHEL 7.6




Figure 1 PowerEdge C4140 Config-M

2.3 Storage components

The storage infrastructure consists of the following components and the details of the configurations can be found from the references:

- Dell EMC Ready Solution for HPC NFS Storage (NSS7.0-HA) (8)
- Dell EMC Ready Solution for HPC Lustre Storage (9)
- Dell EMC Isilon Scale-out NAS Product Family: Dell EMC Isilon F800 All-flash (10)

2.3.1 Dell EMC Ready Solution for HPC Lustre Storage

The Dell EMC Ready Solution for HPC Lustre Storage, referred to as Dell HPC Lustre Storage is designed for academic and industry users who need to deploy a fully-supported, easy-to-use, high-throughput, scale-out and cost-effective parallel file system storage solution. The solution uses Lustre v.2.10.4. It is a scale-out storage solution appliance capable of providing high performance and HA. Using an intelligent, extensive and intuitive management interface, the integrated manager for Lustre (IML), the solution greatly simplifies deploying, managing and monitoring all the hardware and storage system components. It is easy to scale in capacity, performance or both, thereby providing a convenient path to grow in the future. Figure 2 shows the relationship of the metadata server (MDS), metadata target (MDT), management target (MGT), object storage server (OSS) and object storage target (OST) components of a typical Lustre configuration. Further, 2 x PowerEdge R740s are used as MDS. This MDS pair is configured in an active/active, as in case of distributed namespace (DNE) or active/passive for HA. The MDS pair is attached to a PowerVault ME4024, a 2U storage array, through 12 Gbps SAS links which host the MDTs. The OSS pair is attached to 4 x PowerVault ME4048, 5U storage arrays, through 12 Gbps SAS links. The 4x PowerVault ME4048 storage arrays host the OSTs for the Lustre file system. This 4x PowerVault ME4048 configuration is a large configuration, although a medium base configuration is used to benchmark NGS pipelines (11).

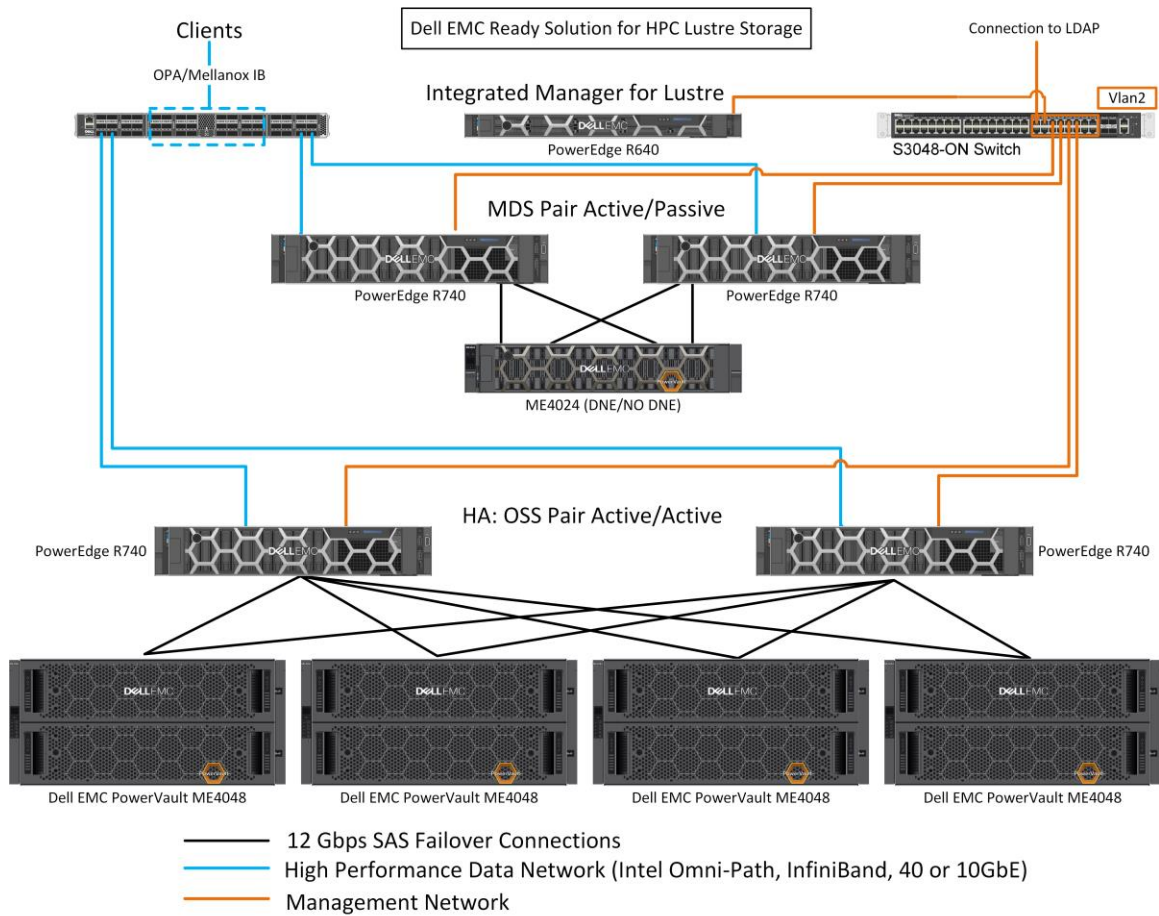


Figure 2 Dell EMC Ready Solution for HPC Lustre Storage, large base configuration including components overview and network configurations.

- Dell EMC PowerEdge R640
 - CPU: 2x Intel Xeon Gold 5118 with 12 cores at 2.3 GHz
 - Memory: 12x 8 GB at 2666 MHz
- 2x Dell EMC PowerEdge R740 for OSS nodes
 - CPU: 2x Intel Xeon™ Gold 6136 CPU with 12 cores at 3.0 GHz
 - Memory: 24x 16 GB at 2666 MHz
 - OSS SAS controller: 4 x Dell 12 Gbps SAS HBAs
 - OSS Storage Array: 4x Dell EMC PowerVault ME4048
 - > Disks: 84-8 TB 7200 RPM NL SAS3 drives per ME4048 enclosure
- 2x Dell EMC PowerEdge R740 for MDS nodes

- CPU: 2x Intel Xeon™ Gold 6136 CPU with 12 cores at 3.0 GHz
 - Memory: 24x 16 GB at 2666 MHz
 - MDS SAS controller: 2 x Dell 12 Gbps SAS HBAs
 - MDS Storage Array: 1 x Dell EMC PowerVault ME4024
- > Disks: 12 x 960 GB SAS SSD (active/passive configuration, no DNE) or 24 x 960 GB SAS SSD (active/Active Configuration, DNE)

2.4 Network components

2.4.1 Dell networking H1048-OPF

Intel Omni-Path Architecture (OPA) is an evolution of the Intel True Scale Fabric Cray Aries interconnect and internal Intel IP (12). In contrast to Intel True Scale Fabric edge switches that support 36 ports of InfiniBand QDR-40Gbps performance, the new Intel Omni-Path fabric edge switches support 48 ports of 100 Gbps performance. The switching latency for True Scale edge switches is 165 ns-175 ns. The switching latency for the 48-port Omni-Path edge switch has been reduced to around 100 ns-110 ns. The Omni-Path host fabric interface (HFI) MPI messaging rate is expected to be around 160 million messages per second (MMPS) with a link bandwidth of 100 Gbps.

2.4.2 Mellanox SB7700

This 36-port Non-blocking Managed InfiniBand EDR 100 Gb/s Switch System provides the highest performing fabric solution in a 1U form factor by delivering up to 7.2 Tb/s of non-blocking bandwidth with 90 ns port-to-port latency.

2.4.3 Dell networking S3048-ON

Management traffic typically communicates with the baseboard management controller (BMC) on the compute nodes using IPMI. The management network is used to push images or packages to the compute nodes from the master nodes and for reporting data from client to the master node. Dell EMC Networking S3048-ON is recommended for management network.

2.5 Reference architecture

This reference architecture is configured with Intel OPA fabric, Dell EMC Ready Solution for HPC NFS Storage and Dell EMC Ready Solution for HPC Lustre Storage. Dell EMC Ready Solution for HPC NFS Storage, referred to as Dell NFS Storage Solution- HA configuration is configured for a general usage such as a home directory, while Dell EMC Ready Solution for HPC Lustre Storage servers as a high-performance scratch storage system. An example solution is illustrated in Figure 3. Although benchmark testing was performed with 64 C6420s, throughput test results showed that 56 compute nodes are optimal with medium base configuration of Dell EMC Ready Solution for HPC Lustre Storage (9).




Figure 3 Dell EMC Ready Solution for HPC Life Sciences with Intel® OPA fabric

Figure 4 illustrates the details of a management network. It requires three Dell Networking S3048-ON switches. Two switches in Rack 1 (Figure 3) make up the internal management network for deploying, provisioning, managing, and monitoring the cluster. One of the switches in Rack 2 has 48 ports that are split into multiple untagged virtual LANs to accommodate multiple internal networks.

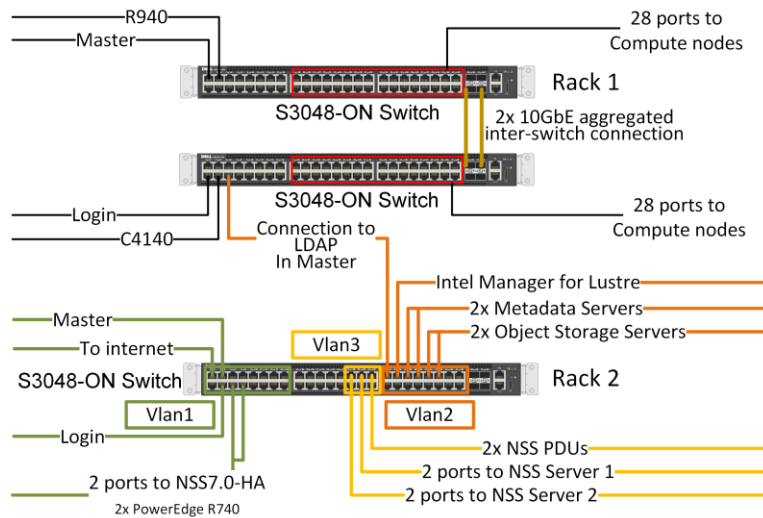


Figure 4 Management network

Figure 5 shows high-speed interconnect with Intel® OPA. The network topology is 2:1 blocking fat tree which requires three Dell Network H1048-OPF 48 port switches.

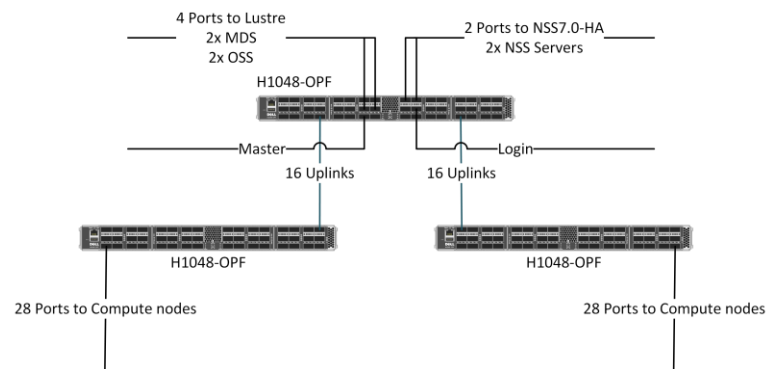


Figure 5 Interconnection with Intel® OPA; 2:1 blocking fat tree topology

2.6 Software components

Along with the hardware components, the solution includes the following software components:

- Bright Cluster Manager
 - > Bright Computing is a commercial application that provides comprehensive software solutions for deploying and managing HPC clusters, big data clusters, and OpenStack in the data center and in the cloud (13). Bright Cluster Manager can be used to deploy complete clusters over bare metal, and manage them effectively. Once the cluster is up and running, the user

interface monitors every single node and reports detection of any software or hardware events.

- BioBuilds

- > BioBuilds is a well maintained, versioned, and continuously growing collection of open-source bio-informatics tools from L7Informatics (14). They are prebuilt and optimized for a variety of platforms and environments. BioBuilds solves most software challenges faced in the life sciences domain.

3 Performance evaluation and analysis

3.1 Variant calling analysis performance

A typical variant calling pipeline consists of three major steps:

- 1) aligning sequence reads to a reference genome sequence;
- 2) identifying regions containing SNPs/InDels; and
- 3) performing preliminary downstream analysis.

In the tested pipeline, BWA 0.7.2-r1039 is used for the alignment step, and Genome Analysis Tool Kit (GATK) is selected for the variant calling step. These are considered standard tools for aligning and variant calling in whole genome or exome sequencing data analysis. The version of GATK for the tests is 3.6, and the actual workflow tested was obtained from the workshop, 'GATK Best Practices and Beyond'. In this workshop, a new workflow with three phases was introduced:

- Best Practices Phase 1: Pre-processing
- Best Practices Phase 2A: Calling germline variants
- Best Practices Phase 2B: Calling somatic variants
- Best Practices Phase 3: Preliminary analysis

Here we tested phase 1, phase 2A and phase 3 for a germline variant calling pipeline. The details of commands used in the benchmark are in APPENDIX A. GRCh37 (Genome Reference Consortium Human build 37) was used as a reference genome sequence, and 50x whole human genome sequencing data from the Illumina platinum genomes project, named ERR194161_1.fastq.gz and ERR194161_2.fastq.gz were used for a baseline test (15).

It is ideal to use non-identical sequence data for each run. However, it is extremely difficult to collect non-identical sequence data having more than 50x depth of coverage from the public domain. Hence, we used a single sequence data set for multiple simultaneous runs. A clear drawback of this practice is that the running time of Phase 2, Step 2 might not reflect the true running time as researchers tend to analyze multiple samples together. Also, this step is known to be less scalable. The running time of this step increases as the number of samples increases. A subtle pitfall is a storage cache effect. Since all the simultaneous runs will read/write roughly at the same time, the run time would be slightly longer in real cases. Despite these built-in inaccuracies, this variant analysis performance test can provide valuable insights when estimating the level of resources required for an identical or similar analysis pipeline with a defined workload.

Total run time is the elapsed wall time from the earliest start of Phase 1, Step 1 to the latest completion of Phase 3, Step 2. Time measurement for each step is from the latest completion time of the previous step to the latest completion time of the current step as illustrated in Figure 6.




Figure 6 Running time measurement method

3.1.1 Single sample test

Three 2nd generation CPUs are picked to compare with four 1st generation CPUs, listed in Table 1. The purpose of this single sample test is to determine suitable CPUs for next generation sequencing (NGS) data analysis. All steps described in Figure 10 are tested on the Dell PowerEdge R640 with 50x whole human genome listed in Table 1 below.

Table 1 Test configuration for single sample variant calling

	Dell PowerEdge R640						
	Intel 1 st Gen Xeon Scalable Processors				Intel 2 nd Gen Xeon Scalable Processors		
CPU	2x 6154	2x 6148	2x 6152	2x 6138	2x 6248	2x 6252	2x 6230
Base Frequency (GHz)	3.0	2.4	2.1	2.0	2.5	2.1	2.1
Number of Cores	18	20	22	20	20	24	20
TDP (W)	200	150	140	140	150	125	125
Memory	24x 16GB DDR4-2666MHz, 2 DPC				12x 32GB DDR4-2933MHz, 1 DPC		
Storage	10x 1.2TB SAS 12 Gbps, 10K in RAID 0						
System Bios	2.1.3						
Kernel	3.10.0-957.el7.x86_64						
OS	Red Hat Enterprise Linux Server release 7.6						
Sequence Reads	ERR194161 , 50x Whole Human Genome for variant calling						

As shown in Figure 7, each step behaves quite differently on each CPU that was tested, and the performance differences among different steps with the tested CPUs ranges from 0.61% to 46.34%. All three 2nd Gen CPUs tested show slightly better performance for most steps in the pipeline. Although 6154 is fastest among tested CPUs, 6154 was not recommended for customers who want to achieve the highest throughput. The fourth step, realigning insertion and deletion step runs on a single core as GATK was not written to utilize multiple cores, and 6248 and 6230 outperformed. Surprisingly, 6230 performs better than 6248 in the fourth step.




Figure 7 Runtime comparisons of steps in variant calling pipeline

Total runtime, sum of all runtime from each step, provides better insight what we can choose for BWA-GATK pipeline as well as other pipelines including aligning and sorting as sub steps. 6248 is the best choice in terms of performance. However, it is important to note that the runtimes in Table 2 are drawn from single sample test (one job on one compute node). The most economical choice could be 6230 due to the relatively low cost and TDP among the tested 2nd Gen CPUs.

Table 2 Total runtime comparisons among 1st and 2nd Gen CPUs

CPU		Price	Spec	Total BWA-GATK runtime (hours)
Skylake	6148	\$3072.00 - \$3078.00	2.4 GHz, 20 cores, 150 W	24.26
	6154	\$3,543.00	3.0 GHz, 18 cores, 200 W	23.47
	6152	\$3655.00 - \$3661.00	2.1 GHz, 22 cores, 140 W	24.58
	6138	\$2612.00 - \$2618.00	2.0 GHz, 20 cores, 125 W	24.83
Cascade Lake	6248	\$3072.00 - \$3078.00	2.5 GHz, 20 cores, 150 W	23.36
	6252	\$3655.00 - \$3662.00	2.1 GHz, 24 cores, 150 W	23.82
	6230	\$1894.00 - \$1900.00	2.1 GHz, 20 cores, 125 W	23.68

3.1.2 Single sample per node, multiple node test

In Figure 8, the runtime in a variety of samples and compute nodes with 50x whole genome sequencing data are summarized. The tests performed here are designed to demonstrate performance at the server level, not to compare individual components. The data points in Figure 8 are calculated based on the total number of samples, one sample per compute node (X axis in the figure) that are processed concurrently. The details of BWA-GATK pipeline information can be obtained from the Broad Institute web site (16). The maximum number of compute nodes used for the tests are 64x C6420s. C6420s with Lustre/ME4 show a better scaling behavior than Lustre/MD3.




Figure 8 Performance comparisons between Lustre/MD3 and Lustre/ME4

3.2 *De novo* assembly performance

For *de novo* assembly, 8168 and 8280M are compared with the same amount of system memory, 1.5 TB in R940. The main reason Cascade Lake 8280M was chosen is for its higher core counts and because it supports more memory, which is beneficial as the data size for *de novo* assembly continues to grow larger over the time.

Table 3 Test bed for *De Novo* Assembly

	Dell PowerEdge R940	
	Intel 1 st Gen Xeon Scalable Processors	Intel 2 nd Gen Xeon Scalable Processors
CPU	4x 8168	4x 8280M
Base Frequency (GHz)	2.7	2.7
Number of Cores	24	28
TDP (W)	205	205
Memory	48x 32 GB DDR4-2666 MHz, 2 DPC	24x 64 GB DDR4-2933 MHz, 1 DPC
Storage	18x 1.2TB SAS 12 Gbps, 10K in RAID 0	
System Bios	2.1.3	
Kernel	3.10.0-957.el7.x86_64	
OS	Red Hat Enterprise Linux Server release 7.6	
Sequence Reads	ERR318658 , 3.2 Billion Reads of Whole Human Genome	

3.2.1 SOAPdenovo2

The maximum performance gain by upgrading from 8168 to 8280M is roughly 1% as shown in 92 cores of 8168 versus 108 cores of 8280M comparisons from Figure 9. For the test, one core per CPU was unused and allocated for the operating system and other housekeeping operations. Although the results show that Cascade Lake 8280M is slower by 2% on average with various number of cores used, the comparisons between 92 cores of 8168 and 108 cores of 8280M confirmed that Cascade Lake 8280M performs slightly better than Skylake 8168.

SOAPdenovo2 seems to be memory bandwidth bounded. The peak memory consumption is constantly rising as more cores are used for a process with the 1 DPC configuration on 2nd Gen CPU, while the peak memory consumption is declining with 2 DPC configuration on 1st Gen CPU. As shown Figure 9 in our previously published blog, memory bandwidth can differ by 11% between 1 DPC and 2 DPC configurations with the same type of dual ranked DIMMs. To come to a more definitive conclusion, further tests are required with the 2 DPC configuration (DDR4-2666) on 2nd Gen 8280M CPU.




Figure 9 Runtime and peak memory consumption plots for SOAPdenovo2 with various number of cores

3.2.2 SPAdes

Cascade 8280M performs better across the tests with various number of cores, and 5% better performance is achievable in a comparison between 92-core 8168 and 108-core 8280M as shown in Figure 10. The patterns of peak memory consumption are nearly similar between the two CPUs; however, 8280M with the 1 DPC configuration shows higher memory consumptions than 8168 with the 2 DPC configuration. Although memory bandwidth does not seem to be critical as we can see from SOAPdenovo2 tests, the 2 DPC configuration with DDR4-2666 MHz can be a better configuration for *de novo* assembly.




Figure 10 Runtime and peak memory consumption plots for SPAdes with various number of cores

4 Conclusion

Overall, 2nd Gen CPUs tested here perform slightly better over Skylake CPUs for genomics workloads such as variant Calling and *de novo* Assembly. In addition to performance advantages, 2nd Gen CPUs provide more choices compared to the predecessors in terms of lower TDP and higher core count for variant calling kinds of workloads. It is notable that 1 DPC configuration with DDR4 2933 MHz DIMMS does not improve performance for SOAPdenovo2. For *de novo* assembly applications, larger memory bandwidth seems to be better; hence, there is no benefit from upgrading memory to DDR4 2933MHz in 1 DPC configuration for 2nd Gen CPUs. It is recommended to set up a 2 DPC configuration with DDR4 2666 MHz, especially for *de novo* assembly applications.

A Technical support and resources

A.1 Related resources

1. Blueprint for High Performance Computing. *Dell TechCenter*. [Online] http://en.community.dell.com/techcenter/blueprints/blueprint_for_hpc/m/mediagallery/20443473.
2. Dell EMC PowerEdge Servers. [Online] <https://www.dell.com/en-us/servers/index.htm>.
3. PowerEdge C6420. [Online] <https://www.dell.com/en-us/work/shop/povw/poweredge-c6420?mkwid=sNuXSS0B5&pcrid=348172952575&pkw=c6420&pmt=e&pdv=c&slid=&product=&pgrid=56791571428&ptaid=aud-637181162098:kwd-378499954786&VEN1=sNuXSS0B5,348172952575,901qz26673,c,,,56791571428,aud-63718116>.
4. Dell EMC Isilon All-Flash. [Online] https://www.dell.com/en-us/storage/isilon/index.htm?cid=68275&st=%2Bemc+%2Bisilon&gclid=CjwKCAjwusrtBRBmEiwAGBPgE5zNeKdC0wUkdc7-_8XWsKx6uv2j7lw_yxLjV8aV8mnT52_i7sROHRoCjEUQAvD_BwE&lid=887928&ptaid=aud-637181162098%3Akwd-308190616027&VEN1=szVFQjcnk%2C32.
5. PowerEdge R940 Rack Server. [Online] https://www.dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r940/pe_r940_12229_vi_vp.
6. PowerEdge C4140 Server. [Online] <https://www.dell.com/en-us/work/shop/povw/poweredge-c4140?mkwid=sApNheUnk&pcrid=348172952803&pkw=c4140&pmt=e&pdv=c&slid=&product=&pgrid=58150336024&ptaid=kwd-402723414919&VEN1=sApNheUnk,348172952803,901qz26673,c,,,58150336024,kwd-402723414919&VEN2=e,c4140>.
7. PowerEdge R740xd Rack Server. [Online] https://www.dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r740xd/pe_r740xd_12238_vi_vp.
8. Dell EMC Ready Bundles for HPC Storage. [Online] <https://si.cdn.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/HPC-Storage.pdf>.
9. Dell EMC Ready Solution for HPC Lustre Storage : Cascade Lake Refresh. [Online] <https://www.dell.com/support/article/us/en/04/sln317174/dell-emc-ready-solution-for-hpc-lustre-storage-cascade-lake-refresh?lang=en>.
10. Dell EMC Isilon F800 AND H600 I/O Performance . [Online] <https://www.emc.com/collateral/white-papers/f800-h600-performance-wp.pdf>.
11. Scalability of Dell EMC Ready Solution for HPC Lustre Storage on PowerVault ME4. [Online] <https://www.dell.com/support/article/us/en/04/sln316413/scalability-of-dell-emc-ready-solution-for-hpc-lustre-storage-on-powervault-me4?lang=en>.
12. Intel True Scale Fabric. [Online] <https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/hpc-true-scale-fabrics.html>.
13. Bright Cluster Manager. [Online] <http://www.brightcomputing.com/products>.
14. BioBuilds. [Online] <https://www.lab7.io/>.

15. European Nucleotide Archive: ERR091571. [Online] <https://www.ebi.ac.uk/ena/data/view/ERR091571>.
16. Genome Analysis Toolkit. [Online] <https://software.broadinstitute.org/gatk/>.