Solution brief

# Ready Architecture for HPC Life Sciences

## Save lives and protect health with more powerful HPC capabilities

### Performance results

**150 50x**

human genomes processed per day[1]

**90 billion**

data points sequenced[2]

**~2X**

throughput improvement over previous system[1]

High Performance Computing (HPC) systems have long been used in life sciences, enabling clinicians and researchers to run complex, compute-intensive workloads. Modern life sciences workloads, such as genomics and proteomics, molecular dynamics, and bioinformatics are defined by large — and growing — data sets that require increasing amounts of computing power.

This exponential growth rate of data requires HPC systems to keep pace, so they can continue enabling researchers to process, analyze and visualize more medical data, faster. The powerful, scalable compute, networking and storage provided by HPC is also necessary to support advanced computing techniques that are trending in life sciences such as artificial intelligence (AI), machine and deep learning.

These advanced computing technologies enable life sciences organizations to keep up with growing workloads, providing the computational horsepower to analyze unprecedented amounts of data. This convergence of HPC, AI and analytics is driving technological advancements that are capable of reshaping the ability of researchers to prevent, detect and treat disease. For example, faster processing speeds enable using AI algorithms that can help make earlier diagnoses, reduce treatment times, accelerate genetic analysis, and speed development of personalized healthcare.

However, designing and deploying an HPC system for life sciences workloads can be challenging. Dell EMC Ready Solutions for HPC Life Sciences are designed to provide faster time to production, better performance and easier scalability to accelerate time to insight for a range of life-sciences use cases. They simplify the configuration, deployment and management of HPC systems. And their modular solution building block approach provides a customizable, flexible architecture that can accommodate various system requirements.

### Dell EMC Ready Solutions for HPC Life Sciences

Dell EMC Ready Solutions for HPC Life Sciences integrates the most relevant of Dell Technologies' expansive portfolio — all tested and tuned for life sciences workloads. The solutions encompass the hardware resources required, while providing an optimal balance of compute density, energy efficiency, and performance.

[1] Dell Technologies Reference Architecture for Dell EMC Ready Solution for HPC Life Sciences, October 2019.

[2] Dell Technologies Case Study, Giving Hope to Children with Rare Disorders, accessed July 2020.

## Resources

- Get validated design and performance testing Information at hpcatdell.com.
- Explore the Dell Technologies HPC & AI Innovation Lab.
- Join the Dell HPC Community at dellhpc.org.

## Learn more

delltechnologies.com/hpc

The Reference Architecture for Dell EMC Ready Solution for HPC Life Sciences outlines the configuration options and performance benchmarks of the latest Dell EMC Ready Solutions for HPC Life Sciences. The paper demonstrates how the solution can process 150 50x human genomes per day using 64x Dell EMC PowerEdge C6420 servers with the Dell EMC Ready Solution for HPC Lustre® Storage. This is roughly a two-fold throughput improvement over the previous version of the solution due to a refresh with the most recent Dell EMC PowerEdge servers, which include the latest Intel® Xeon® Scalable processors with improved memory and storage subsystem performance.[1]

Ready Solutions for HPC Life Sciences have been designed to speed time to production, improve performance with purpose-built solutions, and scale more easily with modular building blocks for capacity and performance. The reference architecture provides guidance for three different configuration options, each optimized for various life sciences applications and workloads such as DNA-Seq, de novo assembly and molecular dynamics simulations. It also compares performance between the new configurations and the previous generation.

The options below serve as a starting point for a customized yet validated HPC life sciences solution.

| PowerEdge Compute nodes (3 options) | PowerEdge Management nodes | Networking | Software and operating system | Storage |
|---|---|---|---|---|
| Choice of:<br>• PowerEdge C6400 enclosure with 4x C6420 Servers<br>• C4140 Server with 4x NVIDIA® V100 PCIe or SXM2<br>• R940 Server | • Master node: R440 Server<br>• Login node and CIFS gateway (optional): R640 Server | • Dell EMC PowerSwitch S3048-ON<br>• Choice of Intel Omni-Path or Mellanox® InfiniBand® EDR 100 Gb/with or without 10/40 GbE | • Bright Cluster Manager®<br>• BioBuilds™<br>• Red Hat® Enterprise Linux® 7.6 | • Dell EMC Ready Solutions for HPC NFS Storage<br>• Dell EMC Ready Solutions for HPC Lustre Storage<br>• Dell EMC Isilon F800 All-flash |

To speed deployment and reduce risk, Dell Technologies HPC experts are available to help you design a solution for your specific needs. And Dell Technologies Services — ranging from consulting and education to deployment and support — are available when and where you need them.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries.

Other trademarks may be the property of their respective owners. Published in the USA 08/20 Solution brief DELL-SB-HPC-LIFESCI-RA-101

Intel® and Xeon® are registered trademarks of Intel Corporation in the U.S. and other countries. Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Bright Cluster Manager® is a trademark of Bright Computing, Inc. Mellanox® and InfiniBand® are registered trademarks of Mellanox Technologies Ltd. Lustre® is a registered trademark of Seagate Technology LLC in the United States.

Dell Technologies believes the information in this document is accurate as of its publication date. The information is subject to change without notice.