Hewlett Packard Enterprise | NVIDIA

Brochure

Supercharging Drug Discovery and Genomics Workflows

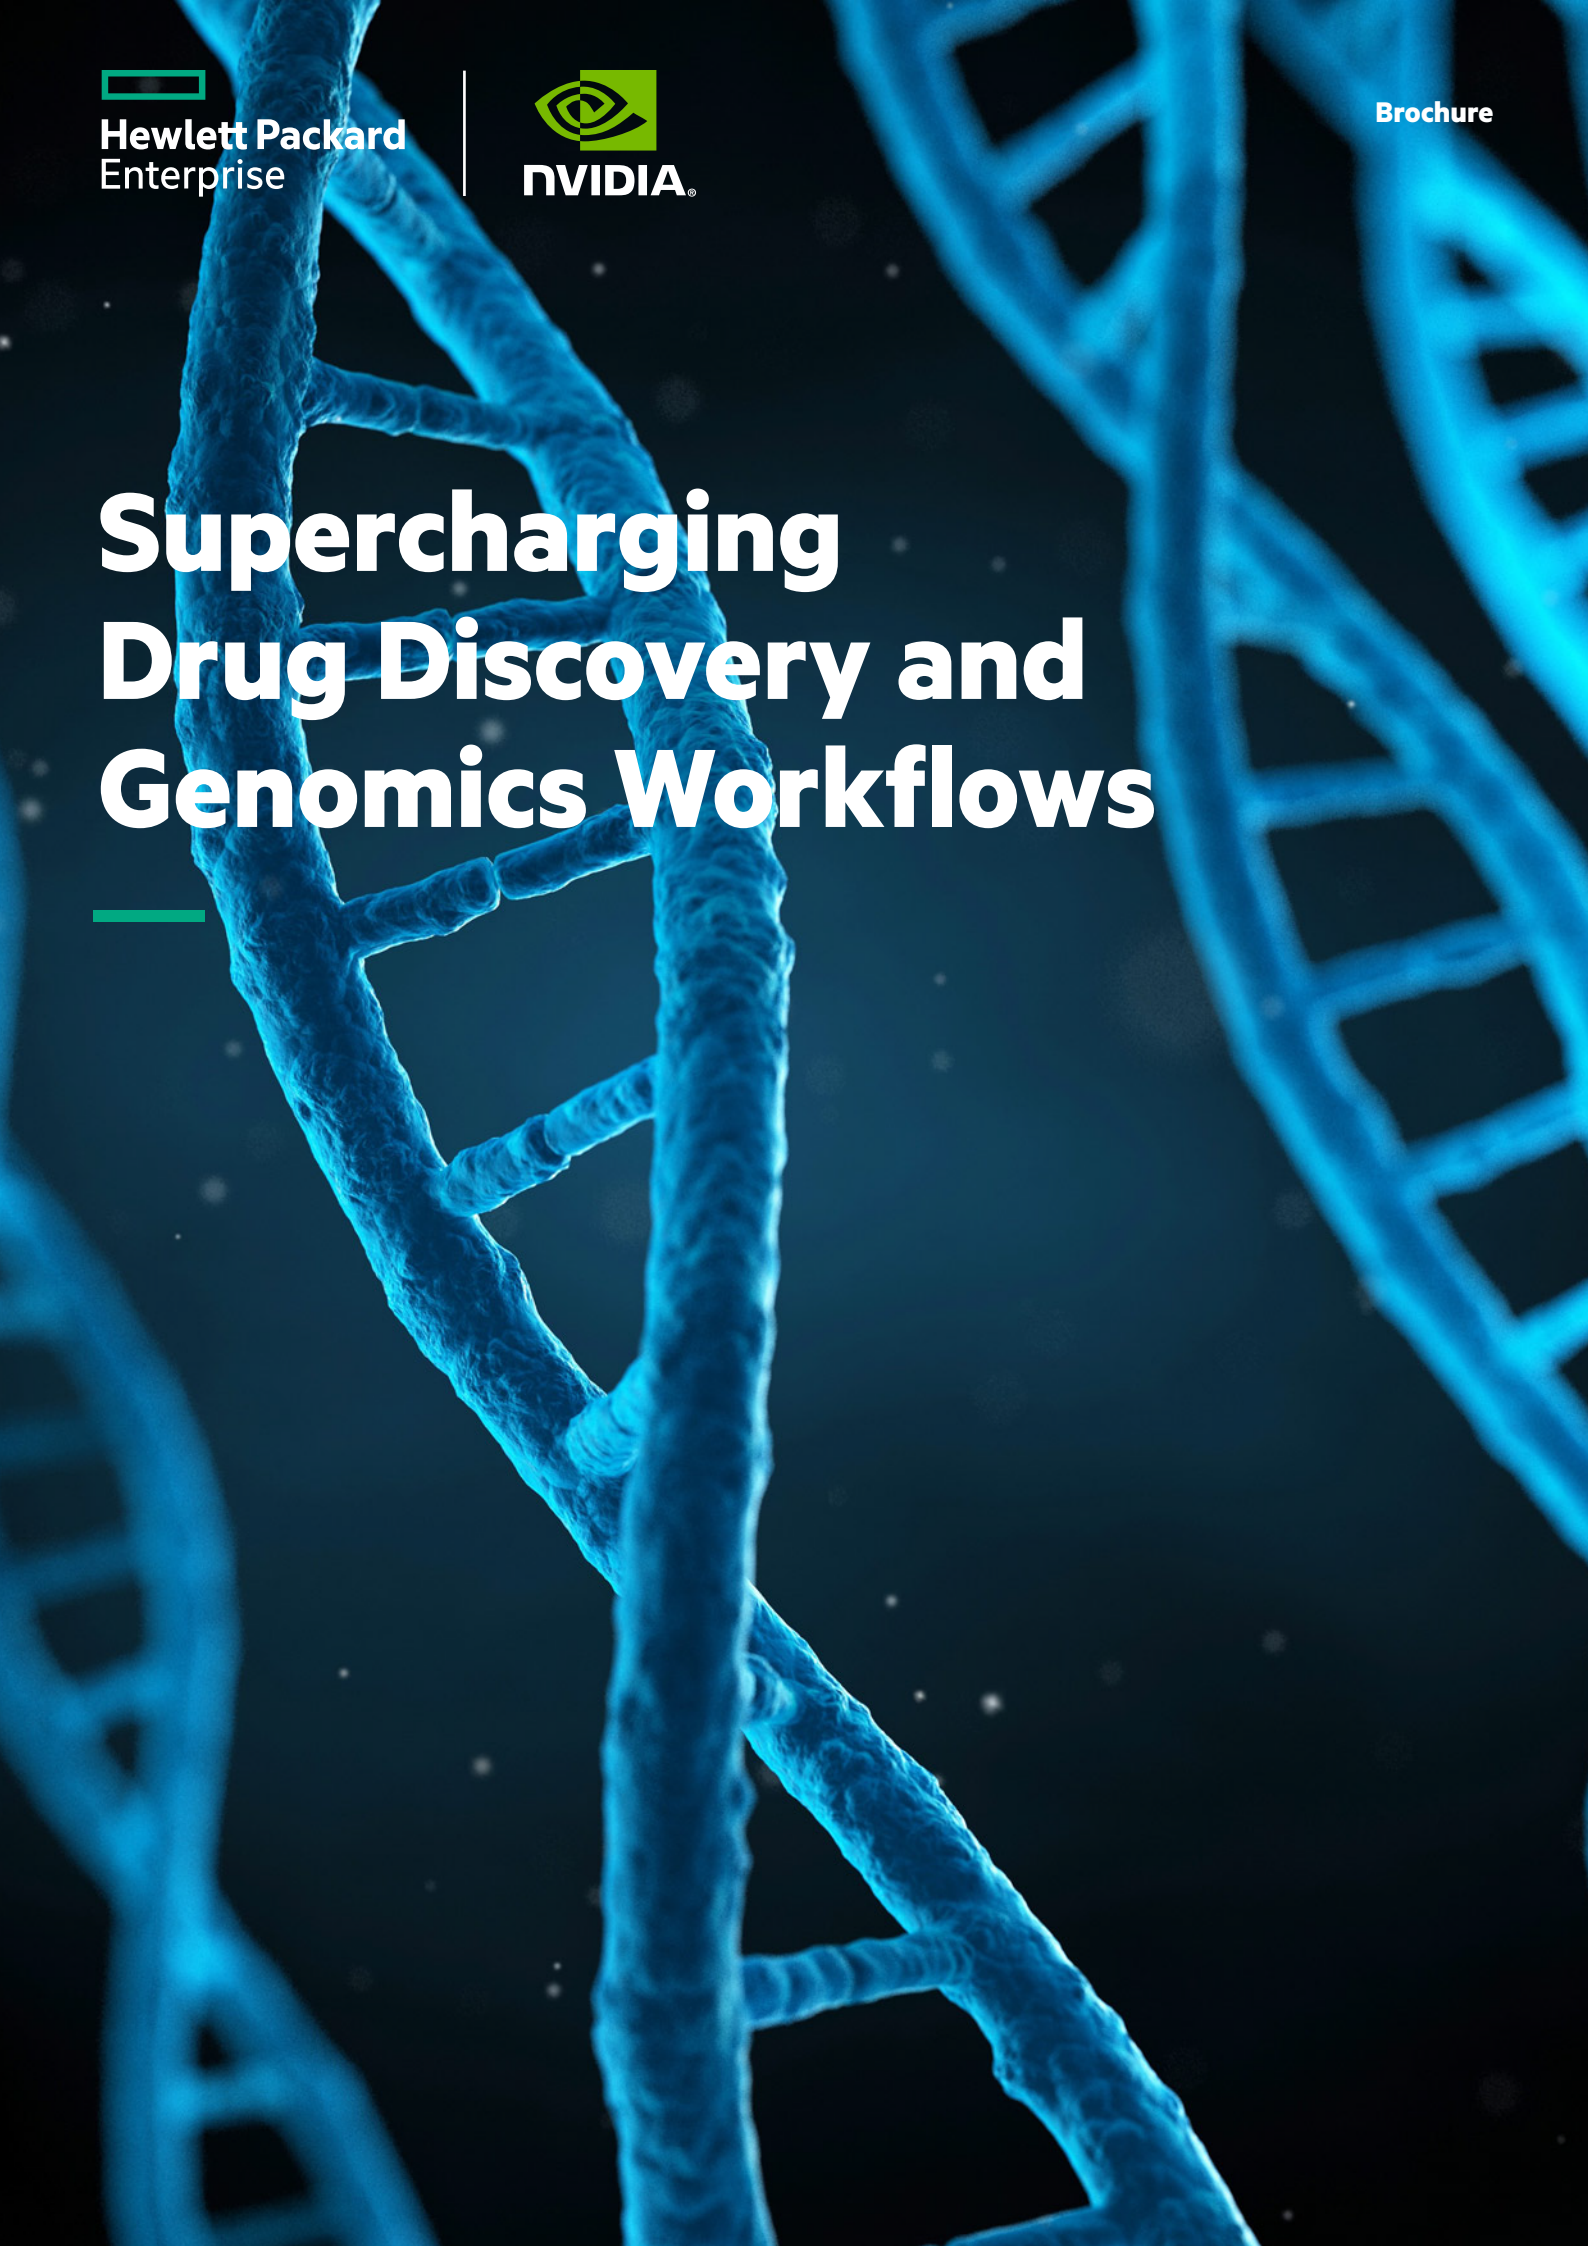

Check if the document is available in the language of your choice.

A new age of life sciences

Artificial intelligence (AI) is transforming the future of drug discovery. Organizations across the life sciences depend on AI to accelerate critical workflows and overcome their greatest challenges with extreme speed and accuracy. These capabilities are key for organizations to advance their understanding of disease biology and deliver the next scientific breakthrough.

The life sciences industry is in a state of evolution. Ongoing advancements in computing are facilitating the investigation, assessment, and creation of new drug development techniques that can be done in silico. These developments can significantly accelerate lead discovery and optimization in drug discovery, leading to improved therapeutics for and prediction of human disease. However, escalating requirements for drug discovery coupled with the exponential growth of life sciences data have put tremendous pressure on organizations to accelerate their research. To accomplish this, organizations must invest in reliable and intelligent solutions to boost their productivity for time-sensitive insight and execute workflows at scale.

Accelerated computing with AI has introduced a revolutionary breed of life sciences applications. Today, organizations are incorporating AI to enhance every phase of drug discovery. The use of AI complements existing workflows by accelerating key processes and advancing research, while providing the same highly accurate results as expensive computational methods. Using AI methods to analyze troves of life sciences data and use insights obtained to accelerate identification of drug candidates and speed time to clinical trial.

HPE and NVIDIA are transforming drug discovery with an AI platform for life sciences. This powerful joint solution is expertly designed to help you achieve better results. The AI platform is the ideal foundation for a new age of life sciences that will empower drug discovery efforts and reveal life-saving discoveries. Organizations in multiple disciplines are implementing the groundbreaking tools for high performance, flexible capacity, secure availability, and simplified management to empower how they work and collaborate. Now, it's your turn.

Digital transformation has become indispensable to life sciences operations. Understanding the potential benefits that AI can bring to your organization begins with understanding how AI streamlines research and development, and why the most demanding applications are AI-powered.

TABLE OF CONTENTS

BEST PRACTICE 1

UNDERSTANDING THE LANDSCAPE OF LIFE SCIENCES

Accelerated computing with AI is solving the greatest challenges in the life sciences industry.

Drug discovery is a lengthy, expensive process, and to a certain extent, it is a shot in the dark. Only one out of thousands of compounds make it through the research, clinical trial, and regulatory review pipeline to be approved and put into large-scale manufacturing. Drugs are now being developed to target specific genes or variants that cause a wide range of diseases. However, the chemical space is vast and researchers face major challenges in searching this space of potential therapeutics for the few candidates that will be ideal for a given target.

Exponential growth of life sciences data has put pressure on organizations to evolve. By 2025, data volumes are expected to reach one zettabase per year. Although this data holds the potential to enable the next scientific breakthrough, issues including computational bottlenecks, limited access to mission-critical information, and poor predictive validity are major roadblocks on the path to discovery. Despite the continued advances in the life sciences space, drug discovery continues to be a struggle—and the problem is increasingly relevant as the ROI for drug discovery R&D is on a downward trend.

Many organizations lack the tools and technologies to perform computationally intensive life sciences data analysis. Without the right infrastructure, it's impossible to run your workflows quickly and efficiently. Parallel computing streams cause delays in processing times due to split and merge approaches. For example, in the genomics workflow, variant reads can cause load imbalances which result in lost productivity and delays. These issues create a large gap between the amount of data that is available and the ability to utilize it effectively. Another challenge is that datasets are often in silos and split between teams and projects.

With legacy technologies running data analysis, the industry has experienced an increase in research and development (R&D) costs. R&D spending is expected to rise from $83 billion in 2019 to roughly $203.9 billion by 2024. Despite this spending, the industry is projected to see a decline in the number of new drugs brought to market. Bringing new drugs to market takes roughly twelve years from drug discovery, pre-clinical studies, clinical trials, and FDA review to finally, large-scale manufacturing. Only one out of thousands of compounds initially identified will become an FDA-approved drug that goes into distribution.

At the same time, life sciences data analysis requirements are growing at an unprecedented rate. For example, there are over two million genomes that will need to be processed in genomics workflow alone by 2025, making genomics poised to be one of the most demanding data domains in existence. With many human diseases in need of cures or treatments, organizations must adopt technologies that will help meet these challenges head on.

BEST PRACTICE 2

BRINGING AI TO POWER EVERY WORKFLOW

AI makes the drug discovery process faster and much more cost-effective, enabling organizations to understand the link between genes, proteins and diseases and to ultimately find cures.

Today, life sciences organizations are unlocking more value from their data with AI. Cutting-edge technologies can dramatically shorten the rigorous and time-consuming process of finding, testing, and optimizing new drugs by helping you identify drug targets and drug candidates faster and also to optimize identified leads more efficiently. AI is helping to make this a predictive process with greater intelligence, higher productivity, and lower costs across all key workflows.

The first stage in the process is lead identification. The more potential compounds an organization can screen, the higher likelihood of identifying ideal drug candidates. It takes an enormous amount of lab time to empirically validate which drug candidate will be optimal, and even more time to optimize selected candidates for efficacy and safety. AI accelerates virtually every computational chemistry code to speed up and reduce the costs of lead identification and optimization and increase the speed with which clinical trials can be started.

Computational chemistry leverages data analytics and in silico simulation to identify and optimize drug targets and lead drug candidates. Computational methods enable you to rapidly screen a large compound library and determine potential binders through modeling, simulation, and visualization techniques. Using cutting-edge analytical instruments such as cryo-electron microscopy (cryo-EM), computational methods that use AI can simulate drug target and lead compound interactions and use structure-based drug design to create better therapeutics. Organizations are reaping the benefits of the latest cryo-EM technology to design and optimize lead compounds in less time and target a wider range of debilitating diseases.

Next generation sequencing (NGS) has revolutionized genomics and fueled the evolution of drug discovery and treatment efforts worldwide. Due to the rapid growth of sequencing workloads, NGS requires innovative AI tools to meet the increasing demand for processing and analysis. The integration of these capabilities drives remarkable improvements in drug discovery by, for example, determining the relationship between genotypes. Many of today's researchers are modeling the human genome against infectious diseases to develop tailored drugs and therapies. Academic research institutions are analyzing petabytes of life sciences data to pinpoint the differences in people's genes, environments, and lifestyles to aid in disease prevention.

Through the use of AI, your organization can meet the demands of drug discovery workflows to ramp up research and enable faster realization of information-driven medicines and treatments:

- Scaling up to handle large influxes of data
- Making data available to multiple teams at the same time
- Accelerating the discovery of new therapies for complex diseases
- Shortening time from lead identification to clinical trial
- Rapidly gaining the economic value from labor-intensive drug discovery workflows
- Mitigating the cost of patent expirations with new discoveries

BEST PRACTICE 3

DEPLOYING A ROBUST AND SCALABLE AI PLATFORM

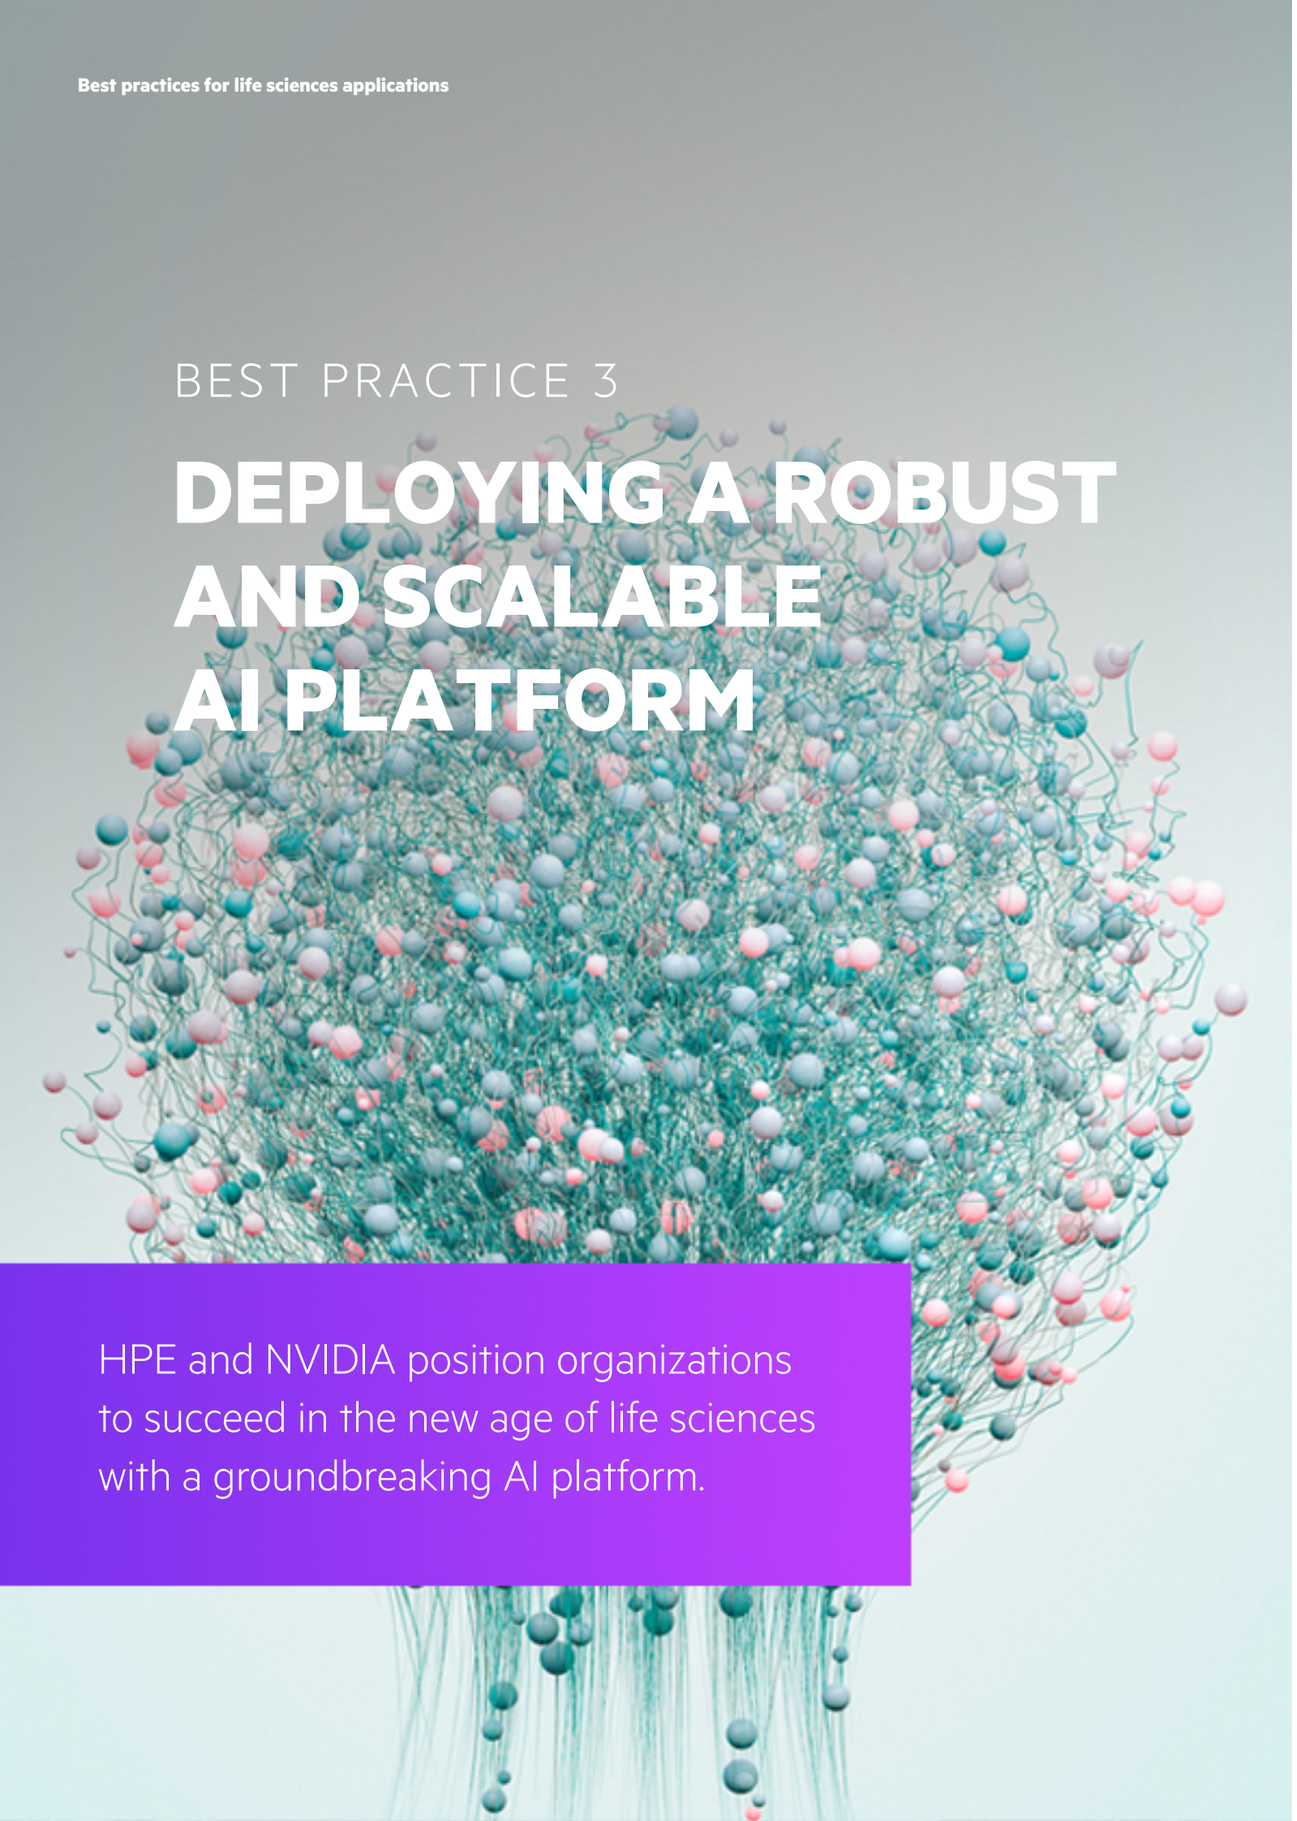

HPE and NVIDIA position organizations to succeed in the new age of life sciences with a groundbreaking AI platform.

To accelerate your drug discovery innovation, you need trusted partners that can deliver breakthrough performance for the most demanding research environments. Hewlett Packard Enterprise and NVIDIA® are advancing the science of drug discovery with better insight, maximum productivity, and faster time to value. We offer a combination of industry-leading technologies that will redefine the possibilities of life sciences. Designed to manage and analyze massive amounts of data, these solutions facilitate end-to-end workflows to accelerate life-changing work.

Supported by a rich independent software vendor (ISV) ecosystem, HPE and NVIDIA offer comprehensive solutions and best practices for reinventing life sciences operations. We work closely with drug discovery ISVs to optimize codes across next generation sequencing, computational chemistry, and structure-based drug design for speed-of-light performance. Organizations around the world leverage these impressive capabilities to overcome data challenges across their disparate teams and projects.

Our robust AI platform is built on HPE systems that are NVIDIA-certified and enable GPU-accelerated applications for supercharging the drug discovery process to maximize ROI. The platform combines compute, storage, interconnects, software, and services for an end-to-end solution. HPE delivers these solutions on-premises, hybrid, or as a service to help simplify system and data management, reduce costs and complexity, and scale to deliver excellent performance. You can choose from an extensive selection of HPE systems that are engineered for AI and powered by NVIDIA GPUs to harness unparalleled processing at any scale.

In addition to vast performance gains over CPU-based platforms, GPU-based HPE systems come with NVIDIA GPUs as well as NVIDIA software tools to optimize workflows that help with AI model creation, AI workflows, and accelerated workflows in drug discovery. NVIDIA AI Enterprise is an end-to-end, cloud-native suite of AI and analytics software that is optimized, certified, and supported by NVIDIA to run on VMware vSphere® with NVIDIA-certified systems. It includes key enabling technologies and software from NVIDIA to manage, secure, and scale AI workloads as well as scale out to multiple nodes for large deep learning training models. Data scientists can run AI workloads at near bare-metal performance to reduce time to deployment of new AI applications. The software delivers high sensitivity and prevision for drug discovery, so you can achieve results that are 99.9% accurate. For GPU-powered NVIDIA Clara Parabricks Pipelines for genomic analysis of DNA and RNA, turnaround times greatly outpace CPU-based workflows with similar accuracy.

NVIDIA Clara Discovery is designed to support cross-disciplinary workflows with the work of thousands of NVIDIA engineering hours behind it. The software brings AI to every stage of the drug discovery process by accelerating computational discovery applications with support for research in computational chemistry, genomics, proteomics, microscopy, virtual screening, visualization, and natural language processing.

NVIDIA Clara Discovery supports drug discovery workflows with NVIDIA Clara Parabricks Pipelines, a core turnkey software component of the AI platform. NVIDIA Clara Parabricks Pipelines significantly accelerates life sciences data analysis on-premises or in the cloud. The software significantly speeds up somatic, germline, and structural variant callers as well as supports mapping and alignment for analysis pipelines. Event-driven insights require multiple threads of execution at once, making them incredibly precise and time-sensitive. For many applications, the window of opportunity for insight is measured in microseconds. The decision loop must complete data capture, analysis, and implement an action immediately. If the process takes too long, the insight is no longer viable.

NVIDIA Clara Parabricks, powered by GPUs, rapidly increases life sciences analysis for somatic, germline, and structural variant workflows for DNA and RNA compared to CPU-based bioinformatic pipelines. The entire pipeline is run using one compute node and does not incur any overhead of distributing data, orchestrating workflows, or reducing accuracy. In fact, the total execution time of the GATK best practices can be dramatically reduced by 59X using NVIDIA Clara Parabricks Pipelines software on a GPU-accelerated HPE system.

Organizations that partner with HPE and NVIDIA are revolutionizing life sciences research with better performance and faster results:

- Screening two billion compounds 33X faster (one day versus three months)
- Screening 10 million drugs 100X faster (eight minutes versus 240 days)
- 76X improvement in throughput performance for whole genome sequencing

BEST PRACTICE 4

HARNESSING AI SOLUTIONS FOR LIFE SCIENCES INNOVATION

HPE makes it simple and cost-effective to deploy an AI platform to empower any environment, so organizations can focus on achieving new life sciences insights.

HPE Pointnext Services offer a broad spectrum of professional, technical, and advisory support for drug discovery—from services like application tuning to more integrated advisory offerings such as project management, on-site consulting, and solution architecture consulting. Our skilled consultants can assist you with installation, configuration, and understanding the requirements of your entire AI deployment.

HPE Pointnext Services make it easy to transform, allowing you to focus on accelerating the performance and impact of data-intensive drug discovery workflows, such as NGS, computational chemistry, and structure-based drug design. Our experts collaborate with you to plan and build your ideal life sciences solutions. We determine your specific goals, application requirements, and potential roadblocks and will even run your technology environment for you after deployment. Several financing and sourcing options are available through HPE Pointnext Services to help you build your solutions and upgrade them on demand.

HPE provides flexible deployment options, so life sciences research institutions, pharma, and biotech companies can choose how they utilize our landmark solutions.

HPE GreenLake cloud services give organizations choice in how they utilize AI. We provide the agility, and scalability of the cloud with the security, simplicity, and control of on-premises IT. Our cloud service business and delivery model leverages data and applications across your operating environments and delivers them anywhere, from edge to cloud. We offer flexible services for you to get started quickly and scale up or down as needed to handle your biggest life sciences and AI workloads. You benefit from a cloud experience when you need it, which helps you avoid heavy upfront investments and expensive overprovisioning by only paying for what you use. Now, you can easily deploy resources, manage costs, and forecast capacity—all from one intuitive platform. HPE GreenLake can manage the entire platform for you with cloud native, zero trust security. With fully managed cloud services, you can focus on your life sciences goals to achieve better clinical and patient outcomes.

Now, you can align IT economics with your AI objectives while simplifying your IT operations, leveraging unmatched speed and capacity.

CONCLUSION

AI is revolutionizing in silico drug discovery. Enabling deep and immediate insights is the key to bring life-saving drugs and treatments to market faster.

As more organizations implement and scale solutions to facilitate research at the cutting-edge of drug discovery, they will need reliable partners with extensive AI and life sciences expertise. As market leaders in life sciences solutions, HPE and NVIDIA are disrupting the industry and powering some of the most demanding research and production environments in the world. Whether building AI for computational chemistry analysis, training a natural language querying tool, or mapping the relationships between diseases, genes, and drugs, our AI platform gives organizations the tools to accelerate value creation. Now, workloads that once took weeks take days, and workloads that took hours take minutes.

Together, we can help prepare every organization to deliver unprecedented health outcomes through better therapeutics. Our solutions are extremely agile and scalable to achieve the next scientific advancement, and we are committed to empowering your success.

Let us help you pioneer the future of life sciences. Contact us today.

LEARN MORE AT

hpe.com/partners/nvidia
hpe.com/us/en/solutions/artificial-intelligence.html

Make the right purchase decision. Contact our presales specialists.

Chat now (sales)

Call now

Get updates

a00119201enw

Hewlett Packard Enterprise