WHITEPAPER

REALIZE THE MASSIVE POTENTIAL OF YOUR AI INFRASTRUCTURE

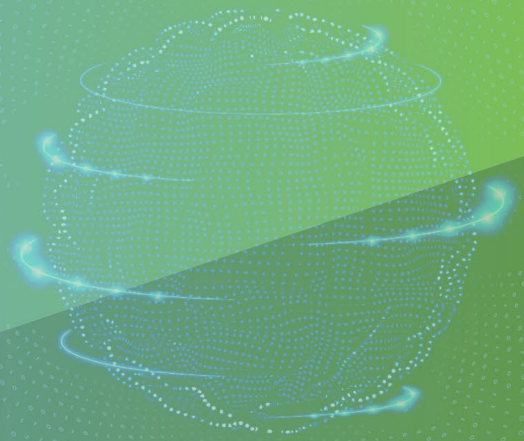

BRIGHT CLUSTER MANAGER ON NVIDIA DGX SYSTEMS

NVIDIA

Bright Computing

REALIZE THE MASSIVE POTENTIAL OF YOUR AI INFRASTRUCTURE

Delivering more of what AI infrastructure users need

NVIDIA DGX servers managed by Bright Cluster Manager deliver the perfect balance of increased concurrency for handling data science workloads, the massive compute requirements associated with those workloads, and the seamless management of all resources across the HPC cluster. With the NVIDIA/Bright solution, HPC users benefit from:

- Solid scale-up architecture for handling data-/compute-intensive workloads
- Easy cluster setup and provisioning
- End-to-end cluster monitoring, health checking, and automated updates
- Automated deployment and configuration of HPC workload managers, Kubernetes, machine learning/deep learning frameworks and libraries, and NGC containers
- Ability to run all the above HPC workloads on the same DGX cluster

INTRODUCTION

Artificial Intelligence (AI) is rapidly becoming an essential business and research tool, providing valuable new insights into corporate data and delivering those insights with high velocity and accuracy. Enterprises, universities, and government organizations are investing tremendous resources to develop a wide array of future-focused Deep Learning (DL) and Machine Learning (ML) solutions such as:

- Autonomous vehicles that circulate unassisted in our cities
- Real-time fraud detection that protects shopping and internet transactions
- Natural language translators that remove language barriers
- Augmented reality that delivers a far richer entertainment experience
- Accelerated drug discovery
- Fully enabled personalized medicine and remote health diagnostics[1]

While these AI capabilities add significant value to our lives, they are the most demanding workloads in modern computing history. The high-performance computing (HPC) clusters required to run AI workloads place significant strain on traditional IT infrastructure, with ever-expanding sets of hot and warm data. Since 2012, the amount of compute used in the largest AI training runs has been increasing exponentially, with a 3.5-month doubling time.[2]

As organizations seek to make AI compute capacity accessible across business units and data science teams, a shared infrastructure approach that is owned and managed by IT (i.e., AI as a Service) is a preferred strategy for enterprises committed to AI development. For a growing number of companies operating in the AI arena, clusters of high-performance NVIDIA® DGX™ servers delivered as a service and managed by Bright Cluster Manager provide the advanced technology and support they need.

[1] DDN Storage, "Scalable Architecture for Artificial Intelligence and Deep Learning with NVIDIA DGX-1," 2018
[2] OpenAI, "AI and Compute," May 16, 2018, https://openai.com/blog/ai-and-compute/

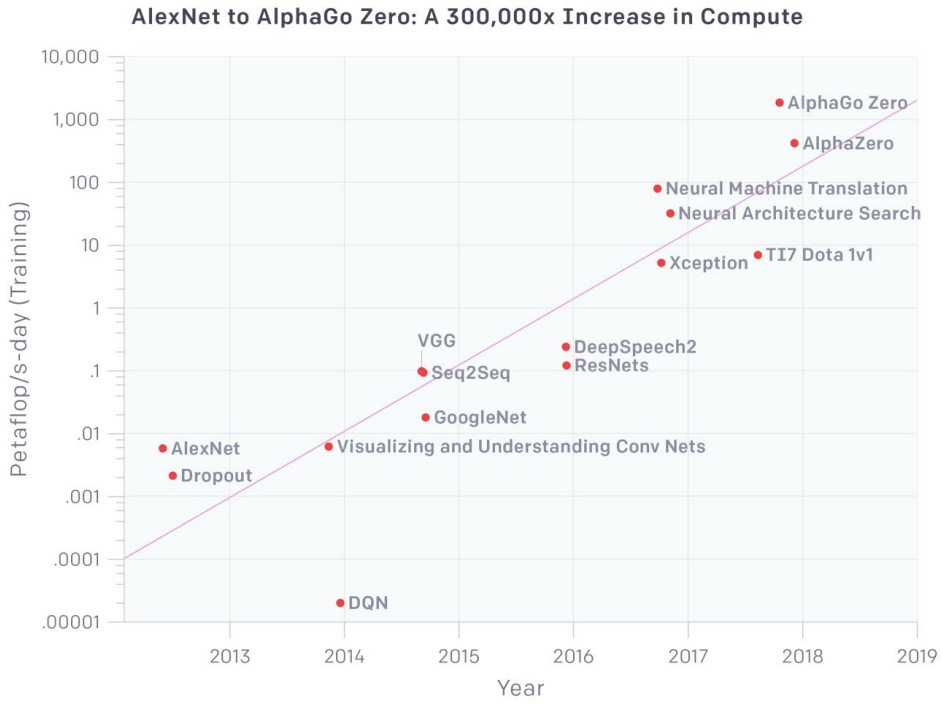

ALL THE VALUE OF AI AND DL, WITH LESS COMPLEXITY AND SEAMLESS CONTROL

Knowing that AI and DL are strategic to the future of nearly all businesses, companies operating in numerous industries—including life sciences/research, media and entertainment, oil and gas, and financial services—want to proactively build AI infrastructure that can accelerate development and speed time to value. Many options exist today, but one that delivers the full potential of AI and DL, with the least complexity and the most seamless control, is NVIDIA DGX with Bright Cluster Manager.

Organizations around the world are familiar with Bright Cluster Manager, trusted for more than a decade to deliver complete and seamless cluster management. NVIDIA DGX is the industry-leading AI system, purpose-built for the unique demands of the enterprise. Working in concert, Bright Cluster Manager and NVIDIA DGX systems deliver an AI infrastructure platform designed to support a wide variety of IT demands—today and over the long term.

With a Bright-managed DGX cluster, organizations not only receive a high-performance platform that enables them to capitalize on AI, but they also benefit from:

- A single tool and interface from which all management and monitoring functions can be performed
- Keeping the compute nodes in sync with each other and with their software image
- Automation of management and monitoring tasks

PUTTING NVIDIA/BRIGHT CLUSTERS TO WORK

Purpose-built for AI, NVIDIA DGX clusters managed by Bright Cluster Manager and enhanced with Bright's Data Science Add-on are the smart choice for a wide variety of end-user use cases. Regardless of use case or workload, Bright Cluster Manager ensures the NVIDIA DGX hardware is utilized efficiently—enabling better performance, ensured service levels, maximum utilization, and improved productivity. Bright also works well with NVIDIA GPU Cloud, which hosts a diverse library of containers performance-tuned for NVIDIA DGX servers, as well as pre-trained models and workflows. While a large number of areas can benefit from DGX and Bright, the most commonly requested areas of help from IT include the following.

Interactive notebooks. When developing a new AI model, an iterative, interactive approach often works best. Instead of submitting jobs that run to completion and waiting to see the results, a Jupyter Notebook (formerly known as IPython) allows users to write one line of code at a time, executing commands as they are written, making it possible to see results in real time; this is a very speedy way to prototype. These interactive notebooks are accessed through a convenient web-based interface, rather than having to log into the command line interface to submit pods or jobs to the cluster using obscure commands. As a platform-agnostic application, Jupyter Notebooks can run on NVIDIA DGX hardware and be managed by Bright Cluster Manager with no customization required. Bright Cluster Manager supercharges Jupyter notebooks at scale by connecting JupyterHub with Kubernetes, Spark, or an HPC scheduler, providing a streamlined way for developers to more easily perform data cleansing and transformation, numerical simulation, statistical modeling, data visualization, machine learning, and more. Users simply request the resources they need for the task at hand, and Bright Cluster Manager delivers the resources to the appropriate Jupyter Notebook through Kubernetes or an HPC scheduler.

AI training and workflows. Training a machine learning model requires a data scientist to provide a learning algorithm, as well as training data the model can learn from. To ensure a well-trained ("smart") machine learning model, engineers build on proven frameworks such as MXNet, PyTorch, TensorFlow, or Theano—all of which are included in the Data Science Add-on for Bright Cluster Manager, without the need for containers. To further simplify DL training, Bright Cluster Manager also provides a choice of machine learning libraries to enable fast and easy access to the datasets the model uses to learn. To enhance performance, DL workflows/processing can be spread across an entire NVIDIA DGX cluster. In addition, Bright Cluster Manager can deploy NGC containers, which are optimized for NVIDIA DGX, providing performance improvements over the upstream branches of the same framework. More complex AI training involves piecing together a workflow that consists of different steps or even a complex DAG (directed acyclic graph). For example, a simple workflow might involve: 1) a first step to import data, analyze it, and transform it; 2) a second step to train a model based on the transformed data; 3) a final step to evaluate the trained model. Each step can involve using a different container image, each potentially leveraging an NGC container for maximum performance. Bright simplifies the deployment of HPC schedulers that facilitate these complex workflows, or Kubernetes apps such as Kubeflow, which provide an alternative solution.

Production model serving. When fully trained and ready for the production stage, a model needs to run at optimal performance not some of the time, but all of the time. Bright Cluster Manager ensures that every workload running on the NVIDIA DGX cluster remains production-ready. With Bright Cluster Manager's built-in support for Docker and Kubernetes, administrators can allocate portions of the NVIDIA DGX cluster resources to container-based applications alongside bare metal applications running on other cluster nodes, and dynamically change resource allocations based on workload demand and/or policy. The result is an efficient production environment with high utilization rates, regardless of the workflow. Using NVIDIA TensorRT Inference Server (known as TRTIS, available on NGC), multiple models can run simultaneously on a single GPU, and they can expand to multi-node deployments as demand increases. TRTIS leverages TensorRT, a library that optimizes trained models so they can run even faster on NVIDIA GPUs.

MAKING LIFE EASIER FOR ADMINISTRATORS

In addition to supporting the primary use cases identified above, a Bright-managed DGX cluster provides critical support for the administrators who manage the clusters and infrastructure. Administrators must increasingly juggle providing more services to their cluster under the pressure and scrutiny of cutting costs.[3] Bright Cluster Manager enables administration of entire clusters as a single entity, provisioning the hardware, operating systems, and workload managers from a unified interface. This approach makes it easier to build, manage, and maintain a reliable NVIDIA DGX cluster.

Initial deployment/installation. Bright Cluster Manager offers an easy-to-use installation wizard to dramatically simplify the process of building a DGX cluster. With a few clicks, Bright Cluster Manager:

- Sets up server network interfaces, packet filtering, and authentication services such as LDAP
- Images all of the cluster's nodes with a DGX-validated OS of choice (Ubuntu, Red Hat Enterprise Linux, or CentOS), workload manager, libraries, and more
- Sets up monitoring and management facilities on each node
- Runs a series of tests across the entire cluster to ensure everything is working properly

Integrated health-checking and monitoring. With Bright Cluster Manager, all monitoring and management are accomplished by a single, highly efficient daemon process that runs on each node in the DGX cluster. Data collected by the node daemons is captured in a single database and consolidated according to preferences. Because Bright's daemon process is highly extensible, the daemon is continuously updated to collect and store new data and metrics exposed by component manufacturers. New hardware automatically shows up in the management interface as it is provisioned into the cluster by the head node. Bright Cluster Manager fully integrates NVIDIA GPUs and provides alerts and actions for metrics such as GPU temperatures, GPU exclusivity modes, GPU fan speeds, system fan speeds, PSU voltages and currents, system LED states, and GPU ECC statistics.

Bright's integrated monitoring can also spot problems proactively—before they result in outages—thereby reducing the impact of failures in production. In most cases, Bright clusters can be updated, application software can be installed, and security patches can be deployed without shutting down or rebooting the cluster nodes. Exceptions exist, but organizations generally perform routine maintenance on their Bright-managed clusters while continuing to use them.

3 Gartner, "Gartner Identifies the Top 10 Trends Impacting Infrastructure and Operations for 2019", December 4, 2018, https://www.gartner.com/en/newsroom/press-releases/2018-12-04-gartner-identifies-the-top-10-trends-impacting-infras

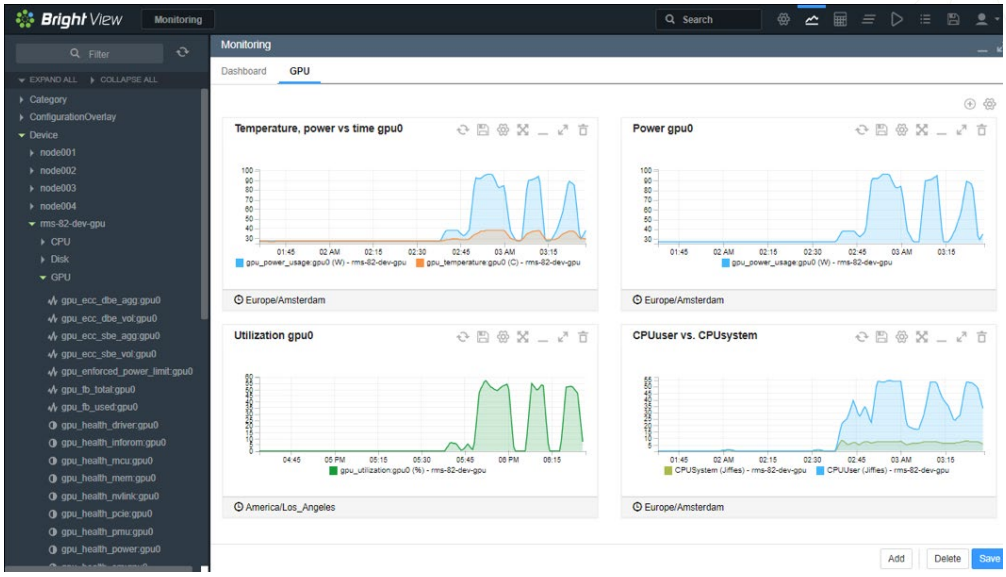

Administrators can visualize monitoring information on a per-device and per-job basis. Bright's workload accounting and reporting functionality can be used to gain insights from monitoring data over longer periods of time (e.g., which users are not using allocated resources efficiently). Bright's health-checking facilities periodically check the health status of resources, such as verifying that parallel filesystems are mounted on nodes, whether user authentication is working properly, and if network interfaces on nodes are functioning correctly.

Maximizing utilization with DGX-as-a-Service. By clustering DGX systems together, administrators can aggregate DGX computing power and offer it "as-a-service" ensuring that demand and supply are aligned and optimized across users, especially when compared to the alternative of independent, siloed DGX nodes spread across the organization. When accessing DGX resources from a cluster, workload engines such as Kubernetes, or HPC schedulers like Slurm, abstract resources and place workloads where they are best serviced across the entire cluster. Bright Cluster Manager installs, configures, and monitors a broad range of workload engines on the DGX cluster simultaneously, allowing users to run different types of applications on a single DGX cluster. By utilizing Bright's workload accounting and reporting capabilities with this as-a-service model, administrators can track DGX resource utilization by user and by job enabling chargebacks and general usage reporting.

Keeping systems up to date. Ensuring the software on cluster nodes remains up to date is critically important; it is also quite burdensome. To alleviate the burden, Bright Cluster Manager automates the process. Administrators can create node categories with different software images to suit specific workloads, and the images can be changed directly from the head node without the need to log in or reboot the affected nodes. Updates to nodes within the cluster can be propagated in one simple step, changes can be easily tracked, and old node images can easily be restored to any node. Packages on node images can be added or removed using standard RPM/dpkg tools or YUM/APT. Administrators can also use popular configuration management tools such as Ansible or Puppet to define the node images that Bright will manage.

NVIDIA DGX SYSTEMS

The world's first portfolio of purpose-built AI systems
Developed to meet the demands of AI and analytics, NVIDIA® DGX™ Systems are built on the revolutionary NVIDIA Volta™ GPU platform. Combined with innovative GPU-optimized software and simplified management tools, these fully-integrated solutions deliver groundbreaking performance and results.

NVIDIA DGX Systems are designed to give data scientists the powerful tools they need for AI exploration—from the desk to the data center to the cloud. With NVIDIA DGX systems, data scientists can experiment faster, train larger models, and arrive at insights—all starting on day one. NVIDIA DGX systems are:

- Powered by the NVIDIA DGX software stack
  - Integrated suite of optimized deep learning software
  - Simplified workload management
- Designed to boost productivity
  - Save hundreds of thousands of dollars in engineering effort
  - Avoid months of the lost productivity spent on IT
- Proven to accelerate return on investment (ROI) for AI
  - Get started in one day instead of weeks or months
  - Accelerate deep learning training by 140X

Comprehensive containerized software stack

NVIDIA DGX Systems deliver an integrated hardware and software solution that has been optimized for delivering the fastest time-to-solution in a GPU-powered appliance. To make this possible, NVIDIA has made significant investment in engineering the deep learning software stack found on every DGX system, as well as the NVIDIA GPU Cloud. This software stack includes the following components:

- **Containerized application**
  - Deep learning applications—including deep learning frameworks, deep learning libraries, and CUDA toolkit
  - NVIDIA GPU access
- **Containerization tool**
  - NVIDIA Container Runtime for Docker
  - Docker engine

GPU-accelerated containers

The NGC container registry provides researchers, data scientists, and developers with simple access to a comprehensive catalog of GPU-accelerated software for AI, machine learning, and HPC. These containers take full advantage of NVIDIA GPUs, both on-premises and in the cloud. Each container is fully optimized and works across a wide variety of NVIDIA GPU platforms and with NVIDIA DGX systems. By delivering predefined containers, NGC containers enable users to focus on building lean models, producing optimal solutions, gathering faster insights, and receiving benefits that matter:

- **Innovate faster**—Get up and running quickly, while reducing the complexity typically associated with setting up software.
- **Remain up to date**—The top DL containers are updated monthly to keep HPC systems running at peak performance; all containers provide easy access to fully tested and optimized software releases.
- **Run anywhere**—NGC containers are built to run on-premises, in the cloud, or in hybrid deployments with Docker and Singularity runtimes; this model allows for maximum utilization of available NVIDIA GPUs, portability, and scalability.

While containerization offers important advantages, deploying containers can be challenging in terms of:

- Provisioning and configuring the underlying servers
- Deploying and managing container orchestration frameworks
- Monitoring the end-to-end infrastructure

NVIDIA DGX with Bright Cluster Manager eliminates these issues and allows organizations to seamlessly run containers alongside other applications within a Bright-managed cluster. The result is an efficient and effective clustered infrastructure with high utilization rates that is ready to support the increasingly diverse needs of users.

Proven technology, exceptional results

Bright Computing has been building enterprise-grade cluster management software for more than a decade. Successfully deployed thousands of times by a thousand organizations around the globe, Bright Cluster Manager is trusted by companies of all sizes operating in nearly all industries. These companies know that Bright Cluster Manager delivers:

- Efficiency—Deploys quickly and easily; enables system-wide monitoring from a single intuitive console
- Availability—Allocates physical and cloud-based resources on demand to meet varying work load requirements and to maintain system availability; repurposes servers to accommodate fluctuating workloads on the fly.
- Extensibility—Extends a physical DGX cluster to the cloud when extra capacity is needed, extends to remote edge locations for latency-sensitive IoT applications, and accommodates a wide range of application types.
- Supportability—Diagnoses and resolves problems and performance issues quickly; deploys software updates across the entire cluster effortlessly.

BRIGHT CLUSTER MANAGER

Advanced cluster management made easy

Bright Cluster Manager eliminates the complexity of installing, managing, and using HPC clusters—making them more capable and flexible without compromising performance. With easy-to-use Bright Cluster Manager, administrators don't need to be Linux™ or HPC experts.

Backed by more than a decade of development—with customer feedback integrated into the product—Bright Cluster Manager embodies the maturity and best practices that can only be gained through real-world use.

With a single-pane-of-glass interface for a cluster's DGX hardware, operating systems, HPC software, networking, and users, Bright Cluster Manager:

- **Deploys easily**—Automatically installs workload managers, container frameworks, deep learning libraries, Spark, and more on the chosen Linux distribution
- **Installs on bare metal**—With nothing to pre-install, enables administrators to quickly build a cluster from bare metal servers or VMs
- **Provides comprehensive monitoring**—Monitors, visualizes, and analyzes a comprehensive set of hardware, software, job- and user-based metrics, and workload accounting and reporting
- **Includes two powerful user interfaces**—Provisions, monitors, and manages clusters via a traditional command line interface or with the web-based graphical user interface called Bright View
- **Integrates hardware management**—Leverages powerful NVIDIA GPU management and monitoring capabilities to provide maximum control of an HPC cluster
- **Optimizes the use of system resources**—Allocates resources according to workload demand or policies to meet business goals
- **Supports containers**—Automatically installs, configures, and manages Docker, Kubernetes, and Singularity across the cluster
- **Includes tools and libraries**—Delivers a complete set of tools and libraries for HPC and deep learning, as well as Apache analytics applications, so administrators can develop, debug, and deploy code immediately
- **Provides comprehensive API**—Allows cluster management operations to be performed from scripts and applications through a JSON-based API; for added convenience, a native Python wrapper is provided around the API, as well as C++ bindings

GETTING THE MOST OUT OF NVIDIA DGX WITH BRIGHT

NVIDIA DGX servers are such powerful compute engines that even a single DGX-1 can benefit from Bright software. One user seldom requires all GPUs on a node, so multiple users are usually allowed to claim GPUs by ID, coordinating usage via calendar, chat, or internally shared spreadsheets. While this method appears to coordinate GPU allocation, some significant problems exist:

- Scales poorly
- Quickly gets outdated
- Is prone to user error

A more elegant solution is to deploy Bright Cluster Manager and use job scheduling, which ultimately results in more efficient use of resources and more productive workers.

Beyond the use of Bright on a single DGX, there are two primary cluster use cases with DGX and Bright Cluster Manager—DGX POD and using DGX servers as part of an existing cluster.

NVIDIA DGX POD

NVIDIA DGX POD™ is an AI infrastructure solution that solves the challenge of scaling enterprise AI— ensuring the optimal balance of compute, storage, and networking to support AI development. Offering an integrated, turnkey approach, DGX POD eliminates design complexity and enables a faster start.

To accelerate and simplify AI deployment, NVIDIA created the DGX POD Ecosystem of IT-Approved Infrastructure Solutions. These solutions leverage technology from numerous industry-leading vendors, including Arista, Cisco, Dell/EMC, DDN, IBM Storage, Mellanox, NetApp, and Pure Storage. Rather than spend time designing an optimal AI infrastructure, IT organizations can use NVIDIA DGX POD solutions to deploy a converged infrastructure in a fraction of the time. Going one step further, all DGX POD-based solutions are backed by 24x7 enterprise-grade support across the entire infrastructure stack—from software to hardware, and from compute to storage to networking

With Bright Cluster Manager, building and managing a DGX POD's clustered compute is fast and efficient. Bright automates a significant portion of the process, as well as simplifies on-going administration of the POD. Once the DGX servers have been racked, powered, and cabled, Bright:

- Images the DGX nodes, including job schedulers
- Sets up and configures networking, firewalls, DNS, LDAP and DHCP
- Sets up monitoring and management facilities on each DGX
- Runs a series of tests across the entire cluster to ensure everything works properly

Using a single Bright interface, the administrator can view all activity on the compute nodes in the POD. In addition, administrators can use the built-in utilities to easily make changes, identify problems, and ensure the system remains up to date and running at peak efficiency.

NVIDIA DGX SERVERS AS PART OF AN EXISTING BRIGHT CLUSTER

As organizations increasingly seek to leverage GPUs for HPC applications and host AI workloads on existing HPC clusters, a growing trend is to add DGX servers to these clusters alongside traditional CPU-based Linux servers. Because Bright Cluster Manager is platform-independent, combining DGX servers with servers from other vendors is easy and seamless. The result—a single unified cluster that Bright can centrally manage and monitor.

When a DGX server is physically added to an existing Bright-managed cluster, the new DGX server is automatically detected, imaged, and configured as a participating node in the cluster. The DGX resources (e.g., GPUs) are made available to users through job schedulers in accordance with the policies defined by the cluster administrator. All monitoring, health-checking, and updating of the DGX are completed through the same interface as all other nodes in the cluster.

For both DGX POD and DGX servers added to existing clusters, the combined power of the DGX system with Bright Cluster Manager translates to:

- **Faster time to value for your AI projects**—NVIDIA's efforts to integrate DGX hardware with its optimized deep learning software stack and create pre-defined AI containers, coupled with Bright's ability to automate the cluster build process and deployment of all software components, results in a DGX cluster being deployed and ready to operate in hours, rather than weeks or months.

- **Bigger AI projects and larger models through effective DGX clustering**—Combining the groundbreaking performance of NVIDIA's DGX with Bright's ability to support clusters of 30,000 or more servers means users can tackle any computational challenge now or in the future; this complementary technology enables HPC environments to scale to support a massive number of end-users.
- **Reduced complexity, reduced administrative burden, and maximum resource utilization**—The technical complexity and administrative burden of compute-intensive AI projects surfaces with the software data scientists need, and the infrastructure upon which that software runs. As the scale of a project grows, so does the complexity of the software and the infrastructure. NVIDIA and Bright have eliminated the complexity of both areas by pre-building, -configuring, and -packaging the software data scientists need, and by automating the process of building, managing, and monitoring the DGX cluster to reduce the time and effort of system administrators.
- **Flexibility and extensibility, from desk to data center to cloud**—In the rapidly unfolding world of AI, the ability to adapt the computing infrastructure is an essential element of success. The adaptive approach must:
    - Support different types of applications, various technologies, numerous geographic locations, and a growing user group
    - Dynamically allocate resources based on business needs and objectives—while also ensuring system resources are properly utilized
    - Easily (automatically) evolve from a silo to a part of the overarching infrastructure

The portability of NVIDIA NGC containers—coupled with Bright Cluster Manager's ability to simultaneously host containers, VMs, and bare metal applications on the same cluster and extend to public clouds—takes the concept of "desk to data center to cloud" to a whole new level.

CONCLUSION

Once used exclusively by elite members of scientific organizations, artificial intelligence has moved to the mainstream—viewed as an essential business and research tool by a wide variety of agencies, companies, and institutions worldwide. As larger and more heterogeneous user groups adopt AI and deep learning technologies, management and maintenance of the AI/DL environment become more complex, costly, and challenging.

A proven solution to AI/DL management is deploying a shared infrastructure that is owned and managed by IT. For a growing number of companies today, a shared infrastructure built on clusters of high-performance NVIDIA DGX servers delivered as a service and managed by Bright Cluster Manager is the solution of choice. Highly flexible, scalable, and extensible, NVIDIA DGX and Bright Cluster Manager deliver:

- More of what AI infrastructure users need
- All the value of AI and DL, with less complexity and seamless control
- Support for a wide range of HPC use cases
- Simplified management
- GPU-accelerated containers

LEARN MORE

Contact your Bright or NVIDIA representative today to learn how your organization can realize the massive potential of an AI infrastructure. You can also visit:

brightcomputing.com
NVIDIA.com/DGX

Tel: +31 20 491 9324
info@BrightComputing.com
www.brightcomputing.com

NVIDIA

Bright Computing