NVIDIA DGX POD for Prescriptive Health and Maintenance of Oil and Gas Systems

RA-09395-001 | March 2019

Reference Architecture

Document Change History

RA-09395-001

Version	Date	Authors	Description of Change
01	2019-03-07	Steven White, Craig Tierney, Ward Eldred, Robert Sohigian	Initial release
02			
03			

Abstract

The NVIDIA DGX POD™ for Prescriptive Health and Maintenance in Oil and Gas (O&G) Platforms reference architecture provides a blueprint for deployment of operational anomaly detection, remaining useful life estimation, and related decision support systems. By leveraging this architecture, it is estimated a 10% improvement in offshore oil production will be achieved by accurately predicting anomalies to reduce unplanned downtimes. This document describes these GPU accelerated deep learning (DL) workloads and provides a guideline for optimizing O&G operations.

A DGX POD includes: one or more racks of NVIDIA® DGX™ servers, storage, networking, NVIDIA GPU Cloud (NGC) DL containers, and DGX POD management servers and software.

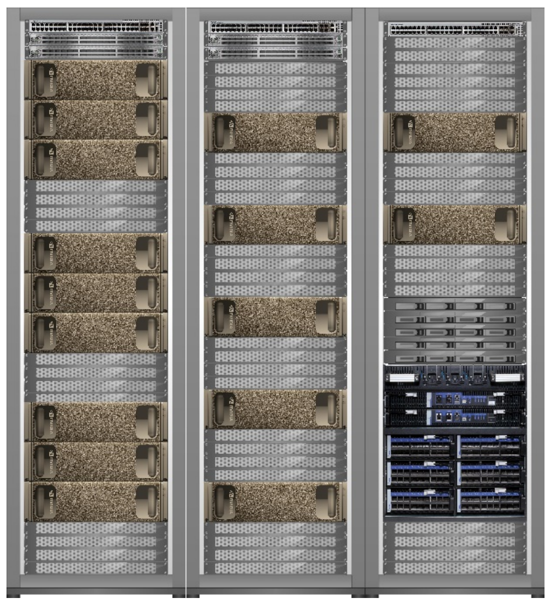

Rack elevations

This reference architecture is based on the NVIDIA DGX SATURNV AI supercomputer which has over 1000 DGX servers and powers internal NVIDIA AI software R&D including medical imaging, autonomous vehicle, robotics, graphics, high performance computing (HPC), and other domains.

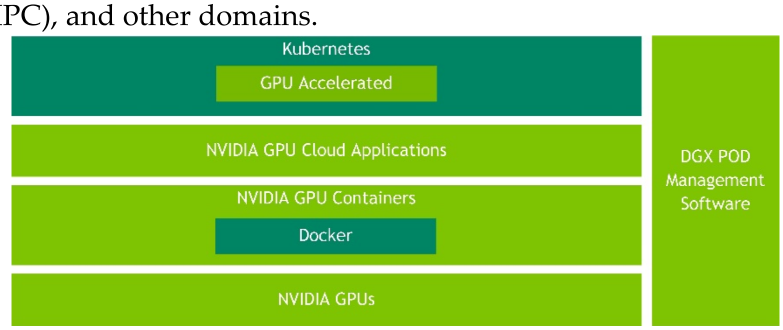

Software stack

Contents

Prescriptive Health and Maintenance Workflow and Sizing

The O&G industry has been collecting large quantities of data for several decades. An offshore oil rig generates 10 TB of data per month on average. However, the industry struggles to unlock end-to-end insights from the data it has been collecting. More than 95% is never used, and hand-crafting features is too slow and doesn't scale well. DL techniques have the potential to significantly improve several aspects of operations in the O&G industry.

Maintaining operational continuity and production efficiency are important focus areas in many industries including O&G. For example, failures of compressors and other equipment are common occurrences in offshore production facilities and refineries. It is estimated that unplanned downtime costs US refiners $6.6 bn in margins each year from seemingly random production disruptions.

Prescriptive health and maintenance (PHM) is a paradigm shift that fosters moving from a dependence on planned events or equipment outages to being able to take real-time action on unexpected events. Usage of DL improves productivity through data-driven insights that has not been possible before with traditional machine learning (ML) techniques. Platforms such as the BHGE AI factory leverage DL and NVIDIA DGX servers for PHM of industrial assets.

A combination of DL and Bayesian algorithms are used to predict the current and future state of industrial assets to diagnose issues and then provide recommendations for improved operations and PHM. Industrial firms are using PHM and analytics software to reduce or eliminate unplanned downtime and maximize profitability and equipment reliability.

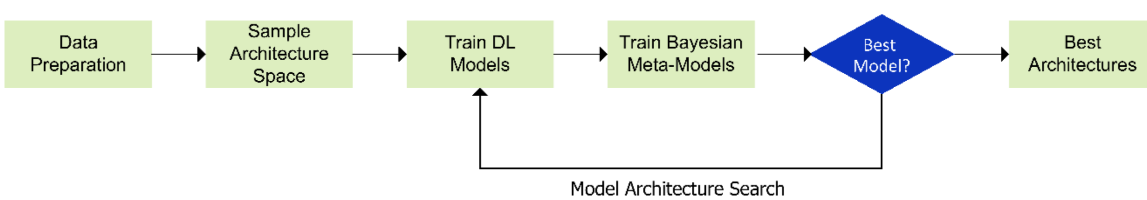

Figure 1. PHM training workflow

A typical PHM training workflow (Figure 1) is as follows:

1. Data preparation. Collects asset sensor data from data sources and validates, transforms, and cleanses the data for model training.
2. Sample architecture space. A pattern search algorithm is trained to automate the identification of source equipment for asset model training.
3. Model Architecture Search. Involves iterating through DL hyperparameters for each model until the optimal network architecture is identified by F1 score.
    a. Autoencoder DL models are trained to learn equipment operational patterns to detect anomalous operational states.
    b. A recurrent neural network (RNN) DL model is trained to detect remaining useful life (RUL) of each asset.
    c. Natural language processing (NLP) models are trained on equipment manuals and service records to provide root cause analysis and prescriptive guidance for failure maintenance.
    d. A Bayesian model is trained to provide a confidence score for each model classification.
4. Best Architectures. The best models for each of the three use cases (RUL, anomaly detection, and prescriptive maintenance) are then deployed for inferencing in production.

A key part of this training workflow is the model architecture search that leverages a genetic search algorithm to iterate through the network architecture hyperparameters such as number of hidden layers, type of activation functions, learning rates, etc., to find the best performing model based upon F1 score. The search space is generally large and the response (i.e. best generalizing model with least error) could be highly non-linear. An ad-hoc approach to finding a good model through manual trial and error and relying on the data scientist’s experience may result in suboptimal models. To consistently and efficiently find the best model architecture, a combination of DL models and probabilistic models are required.

This hyperparameter search is a computationally expensive process that can be significantly accelerated with GPUs. For example, a typical optimization process involves searching through 250,000 model architectures. A new architecture search is generally expected every six months per model. Any significant change in operations of equipment would require an architecture search.

Sizing of DL training and inferencing is highly dependent on data size and model complexity. For example, a predictive maintenance training workflow creates an ensemble of models to include equipment and mode identification, autoencoders for anomaly detection, natural language processing trained on the industry-specific lexicon, and RUL models for detecting equipment failure.

Each of these models are trained using several years of data. This can reach upwards of 100 TB for large offshore platforms and refineries, which have hundreds of assets and sub-components.

Typical sizing parameters for offshore platforms are shown in Table 1.

Use-Case	Systems per Offshore Platform	Industrial Assets per System	DGX-1 Servers	
			Inferencing	Training
RUL	20	20	2+1[1]	13
Anomaly Detection	20	20	1	7
Prescriptive Maintenance	20	20	1	2

1. Due to the mission critical nature and real-time needs of monitoring asset health, inferencing servers should be deployed in an N+1 High Availability configuration.

Table 1. Typical sizing parameters for offshore platforms

Once the models for each of the use cases have been tested and are ready for deployment, the latency, throughput, and efficiency of the AI model is optimized for production deployment (inferencing) using the NVIDIA TensorRT™ inference accelerator.

Each of the use cases are detailed in the following sections.

Predicting Remaining Useful Life

Conventional diagnostics are generally reactive (i.e., provides mitigating actions to limit damage after failure has occurred). Providing a proactive solution by predicting impending failures in advance can lead to reduced downtime and maintenance costs. Accurate time-to-failure predictive DL models can be built by leveraging historical time series sensor data collected from assets in the field.

RNNs are used for predicting remaining useful life because they can handle sequence or time dependencies. Specifically. a long short-term memory (LSTM) network will be used as it is more flexible and stable than standard RNN models.

Table 2 shows the estimate of DGX-1 servers required to support training and inference the RUL workflow. Sizing is based on a generic offshore platform consisting of 400 assets and sub-components with each asset consisting of 100 sensors.

Industrial assets and sub-components	**Sensors**	**Data per Quarter**	**DGX-1 Servers**	
			Inferencing	**Training**
400	40,000	15 TB	2+1[1]	13

1. Due to the mission critical nature and real-time needs of monitoring asset health, inferencing servers should be deployed in an N+1 High Availability configuration.

Table 2. RUL Sizing

Anomaly Detection

Production continuity and production efficiency is an important focus for the O&G industry. For example, failures of compressors and other equipment are common occurrences in offshore production facilities. The random nature of these failures results in production disruption that amounts to hundreds of millions of barrels of lost or deferred oil production annually.

Anomaly detection is a useful technique to apply to time series problems where you have a large imbalance of successful vs failure events, such as industrial equipment that has been engineered with large mean time between failures (MTBF) as well as cases where you have un-labeled data. Specifically, a long short-term memory (LSTM) autoencoder will be used for unsupervised learning of temporal events and signal patterns from input data. A custom-tailored model for each asset or sub-system will be created.

Table 3 shows the estimate of DGX-1 servers required to support training and inference the anomaly detection workflow. Sizing is based on a generic offshore platform consisting of 400 assets and sub-components with each asset consisting of 100 sensors.

Industrial Assets and Sub-Components	Sensors	Data per Quarter	DGX-1 Servers	
			Inferencing	Training
400	40,000	15 TB	1	7

Table 3. Anomaly Detection Sizing

Prescriptive Maintenance

Traditional industries like the O&G industry possesses decades of data that has been recorded and continues to generate additional data every day. This includes everything from field notes, service records to equipment operating manuals and technical reports and publications. This textural data is an invaluable resource for optimizing operations. DL and natural language processing (NLP) have revolutionized language models and can be used to both understand and make recommendation to field engineers trying to resolve issues.

NVIDIA works with industry partners such as BHGE that have created a set of NLP models specifically on O&G lexicon to understand industry specific nuances and acronyms in both spoken and written language. The NLP models can then be used to train on the service records from years of field service data. To provide a baseline for sizing hardware for training NLP models, consider the following outcomes typically required as part of diagnosis and prognosis:

- Questions and answering for understanding potential root causes of failure
- Potential steps for debugging equipment
- Procedures for regular maintenance

The R-Net architecture using 20 million words or roughly 40,000 pages of text was used to create the sizing example shown in Table 4.

Words	Pages	DGX-1 Servers	
		Inferencing	Training
20M	40,000	1	2

Table 4. Prescriptive maintenance sizing

Such a model would need to be re-trained and kept up to date on operational manuals, both from customer and external vendors.

NVIDIA AI Software Stack

NVIDIA AI and HPC software running on the DGX POD provides a high-performance environment for large scale multi-user software development teams. NVIDIA AI software (Figure 2) includes the DGX operating system (DGX OS), cluster management and orchestration tools, NVIDIA libraries and frameworks, workload schedulers, and optimized containers from the NGC container registry. To provide additional functionality, the DGX POD management software includes third-party open-source tools recommended by NVIDIA which have been tested to work on DGX PODs with the NVIDIA software stack. Support for these tools can be obtained directly through third-party support structures.

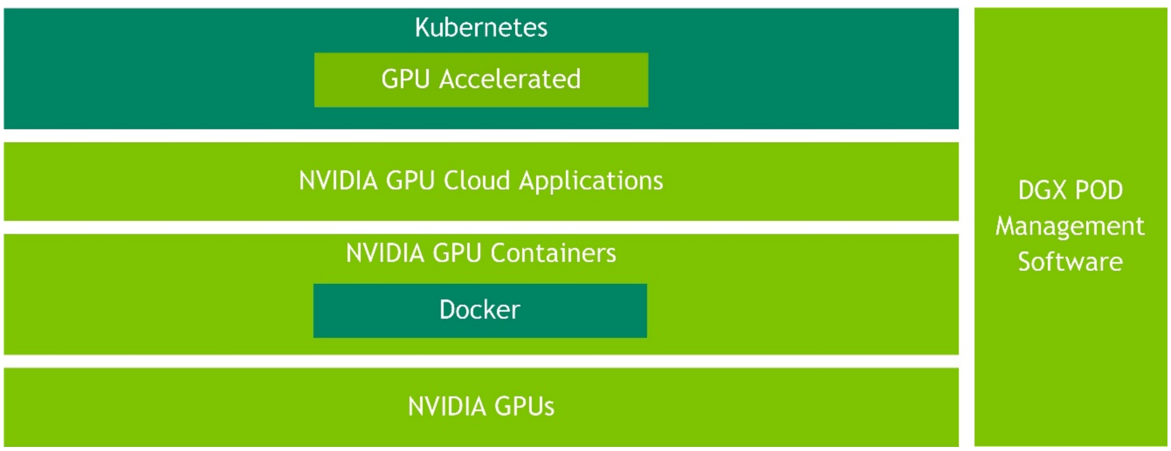

Figure 2. NVIDIA AI software stack

The foundation of the NVIDIA AI software stack is the DGX OS, built on an optimized version of the Ubuntu Linux operating system and tuned specifically for DGX hardware.

The DGX OS software includes certified GPU drivers, a network software stack, pre-configured NFS caching, NVIDIA data center GPU management (DCGM) diagnostic tools, GPU-enabled container runtime, NVIDIA CUDA® SDK, NCCL and other NVIDIA libraries, and support for NVIDIA GPUDirect™ technology. The DGX OS software can be automatically re-installed on demand by the DGX POD management software.

The management software layer of the DGX POD (Figure 3) is composed of various services running on the Kubernetes container orchestration framework for fault tolerance and high availability. Services are provided for network configuration (DHCP) and fully-automated DGX OS software provisioning over the network (PXE).

Monitoring of the DGX POD utilizes Prometheus for server data collection and storage in a time-series database. Cluster-wide alerts are configured with Alertmanager, and DGX POD metrics are displayed using the Grafana web interface. For sites required to operate in an air-gapped environment or needing additional on-premises services, a local container registry mirroring NGC containers, as well as Ubuntu and Python package mirrors, can be run on the Kubernetes management layer to provide services to the cluster.

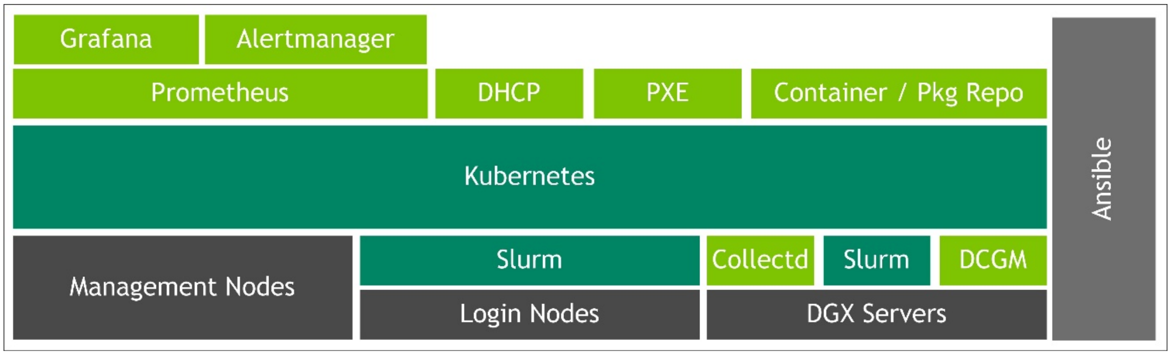

Figure 3. DGX POD management software

The DGX POD software allows for dynamic partitioning between the nodes assigned to Kubernetes and Slurm such that resources can be shifted between the partitions to meet the current workload demand. A simple user interface allows administrators to move DGX-1 servers between Kubernetes- and Slurm-managed domains.

Kubernetes serves the dual role of running management services on management nodes as well as accepting user-defined workloads and is installed on every server in the DGX POD. Slurm runs only user workloads and is installed on the login node as well as the DGX server nodes. The DGX POD software allows individual DGX servers to run jobs in either Kubernetes or Slurm. Kubernetes provides a high level of flexibility in load balancing, node failover, and bursting to external Kubernetes clusters, including NGC public cloud instances. Slurm supports a more static cluster environment but provides advanced HPC-style batch scheduling features including multi-node scheduling that some workgroups may require. With the DGX POD software, idle servers can be moved back and forth as needed between the Kubernetes and Slurm environments. Future enhancements to Kubernetes are expected to support all DGX POD use cases in a pure Kubernetes environment.

User workloads on the DGX POD primarily utilize containers from NGC (Figure 4), which provides researchers and data scientists with easy access to a comprehensive catalog of GPU-optimized software for DL, ML, HPC applications, and HPC visualization that take full advantage of the GPUs. The NGC container registry includes NVIDIA tuned, tested, certified, and maintained containers for top ML libraries and top DL frameworks such as TensorFlow, PyTorch, and MXNet. NGC also has third-party managed HPC application containers, and NVIDIA HPC visualization containers.

Figure 4. NGC overview

Management of the NVIDIA software on the DGX POD is accomplished with the Ansible configuration management tool. Ansible roles are used to install Kubernetes on the management nodes, install additional software on the login and DGX servers, configure user accounts, configure external storage connections, install Kubernetes and Slurm schedulers, as well as performing day-to-day maintenance tasks such as new software installation, software updates, and GPU driver upgrades.

The software management stack and documentation are available as an open source project on GitHub at:

> https://github.com/NVIDIA/deepops

DGX POD Design

The DGX POD is an optimized data center rack containing up to nine DGX-1 servers, storage servers, and networking switches to support single and multi-node AI model training and inference using NVIDIA AI software.

This section details three DGX POD designs:

- O&G POD (35 kW)
- O&G POD (18 kW)
- O&G POD Utility Rack

There are several factors to consider when planning a DGX POD deployment to determine whether to use the 35 kW or 18 kW configurations. These reference architectures are designed to utilize high-density racks to provide the most efficient use of costly data center floorspace and to simplify network cabling. As GPU usage grows, the average power per server and power per rack continues to increase. However, since older data centers may not yet be able to support the power and cooling densities required, the 18 kW design should be used.

These DGX POD elevations are designed to fit within a standard-height 42 RU data center rack. A taller rack can be used to include redundant networking switches, a management switch, and login servers. This reference architecture uses an additional utility rack for login and management servers and has been sized and tested with up to six DGX PODs. Larger configurations of DGX PODs can be defined by an NVIDIA solution architect.

A primary 10 GbE (minimum) network switch is used to connect all servers in the DGX POD and to provide access to a data center network. The DGX POD has been tested with an Arista switch with 48 x 10 GbE ports and 4 x 40 GbE uplinks. VLAN capabilities of the networking hardware are used to allow the out-of-band management network to run independently from the data network, while sharing the same physical hardware. Alternatively, a separate 1 GbE management switch may be used. While not included in the reference architecture, a second 10 GbE network switch can be used for redundancy and high availability. In addition to Arista, NVIDIA is working with other networking vendors who plan to release switch reference designs compatible with the DGX POD.

A 36-port Mellanox 100 Gbps switch is configured to provide four 100 Gbps InfiniBand connections to the DGX servers in the rack. This provides the best possible scalability for multi-node jobs. In the event of switch failure, multi-node jobs can fall back to use the 10 GbE switch for communications. The Mellanox switch can also be configured in 100 GbE mode for organizations that prefer to use Ethernet networking. Alternately, by configuring two 100 Gbps ports per DGX server, the Mellanox switch can also be used by the storage servers.

With the DGX family of servers, AI and HPC workloads are fusing into a unified architecture. For organizations that want to utilize multiple DGX PODs to run cluster-wide jobs, a core InfiniBand switch is configured in the utility rack in conjunction with a second 36-port Mellanox switch.

Storage architecture is important for optimized DL training performance. The DGX POD uses a hierarchical design with multiple levels of cache storage using the DGX SSD and additional cache storage servers in the DGX POD. Long-term storage of raw data can be located on a wide variety of storage devices outside of the DGX POD, either on-premise or in public clouds.

The DGX POD baseline storage architecture consists of in-rack NFS storage servers used in conjunction with the local DGX SSD cache. Additional storage performance may be obtained by using one of the distributed filesystems listed at this link:

https://docs.nvidia.com/deeplearning/dgx/bp-dgx/index.html#storage_parallel

The DGX POD is also designed to be compatible with several third-party storage solutions, many of which are documented at this link:

https://www.nvidia.com/en-us/data-center/dgx-reference-architecture/

O&G POD (35 kW)

Figure 5 shows nine DGX-1 severs in a 35 kW rack configuration. Major components and their rack units are detailed next to the figure.

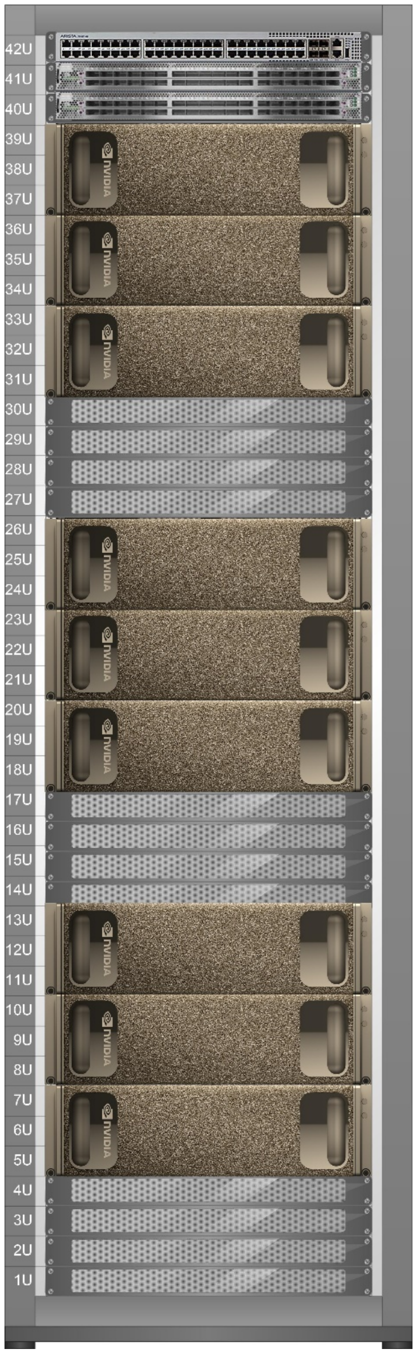

- Nine DGX-1 servers
  (9 x 3 RU = 27 RU)
- Two 100 Gbps intra-rack network switches
  (2 x 1 RU = 2 RU)
- One 10 GbE (min) storage and management switch
  (1 RU)

Figure 5. O&G 35 kW elevation

O&G POD (18 kW)

Figure 6 shows five DGX-1 severs in an 18 kW rack configuration. Major components and their rack units are detailed next to the figure.

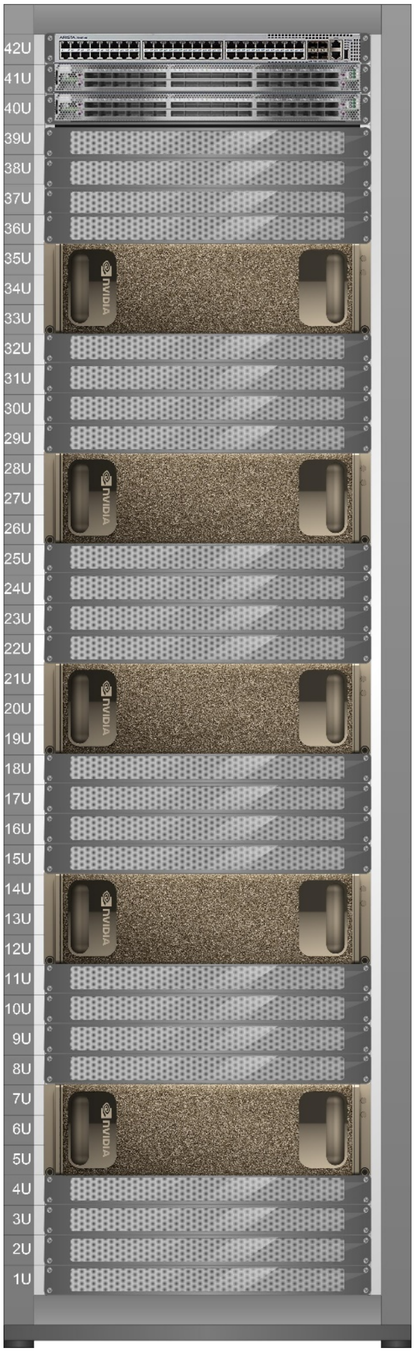

- Five DGX-1 servers (5 x 3 RU = 15 RU)
- Two 100 Gbps intra-rack network switches (2 x 1 RU = 2 RU)
- One 10 GbE (min) storage and management switch (1 RU)

Figure 6. O&G 18 kW rack elevation

O&G POD Utility Rack

Figure 7 shows a utility rack for login, storage, and management servers as well as management and clustering switches. Major components are detailed next to the figure.

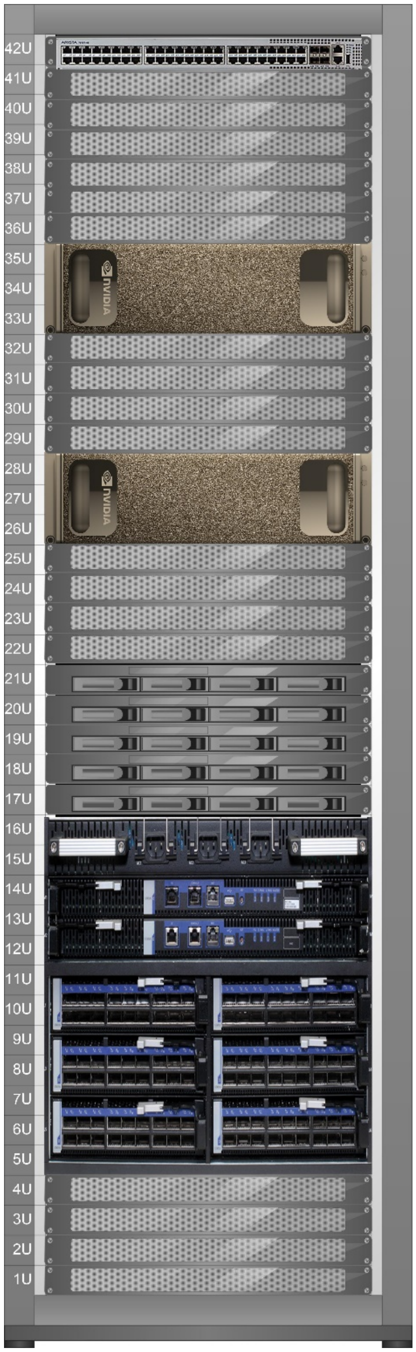

- Two DGX-1 servers (needed for 18 kW configuration)

  (2 x 3 RU = 6 RUs)
- Login server which allows users to login to the cluster and launch Slurm batch jobs[1].

  (1 RU)
- Three management servers running Kubernetes server components and other DGX POD management software[2].

  (3 x 1 RU = 3 RUs)
- One storage server (ten 1 TB disk drives)

  (1 RU)
- Optional multi-POD clustering using a Mellanox 216 port EDR InfiniBand switch.

  (12 RUs)
- Optional multi-POD 10 GbE storage and management network switches (not shown)

  (2 x 1 RU = 2 RUs)

[1] To support many users, the login server should have two high-end CPUs, at least 1 TB of memory, two links to the 100 Gbps network, and redundant fans and power supplies.

[2] These servers can be lower performance than the login server and can be configured with mid-range CPUs and less memory (128 to 256 GB).

Figure 7. O&G utility rack elevation

DGX POD Installation and Management

DGX POD deployment is like deploying traditional servers and networking in a rack However, with high-power consumption and corresponding cooling needs, server weight, and multiple networking cables per server, additional care and preparation is needed for a successful deployment. As with all IT equipment installation, it is important to work with the data center facilities team to ensure the DGX POD environmental requirements can be met.

Additional DGX server site requirements are detailed in the *NVIDIA DGX Site Preparation Guide* but important items to consider include:

<table>
<tr><th rowspan="2">Area</th><th colspan="2">Design Guidelines</th></tr>
<tr><th>O&G POD (35 kW)</th><th>O&G POD (18 kW)</th></tr>
<tr><td rowspan="2">Rack</td><td>Supports 3000 lbs. of static load</td><td>Supports 1200 lbs. of static load</td></tr>
<tr><td colspan="2">• Dimensions of 1200 mm depth x 700 mm width
• Structured cabling pathways per TIA 942 standard</td></tr>
<tr><td rowspan="2">Cooling[1]</td><td>Removal of 119,420 BTU/hr</td><td>Removal of 59,030 BTU/hr</td></tr>
<tr><td colspan="2">ASHRAE TC 9.9 2015 Thermal Guidelines “Allowable Range”</td></tr>
<tr><td>Power</td><td>• North America: A/B power feeds, each three-phase 400V/60A/33.2kW (or three-phase 208V/60A/17.3 kW with additional considerations for redundancy as required)
• International: A/B power feeds, each 380/400/415V, 32A, three-phase –21-23kW each.</td><td>• North America: A/B power feeds, each three-phase 208V/60A/17.3 kW
• International: A/B power feeds, each 380/400/415V, 32A, three-phase –21-23kW each.</td></tr>
<tr><td colspan="3">1. Via rack cooling door or data center hot/cold aisle air containment</td></tr>
</table>

Table 5. Rack, cooling, and power considerations for DGX POD racks

Figure 8 shows the server components and networking of the DGX POD. management servers, login servers, DGX servers, and storage communicate over a 1 or 10 Gbps Ethernet network, while login servers, DGX servers and optionally storage can also communicate over high-speed 100 Gbps InfiniBand or Ethernet. The DGX servers shown here are running both Kubernetes and Slurm to handle varying user workloads.

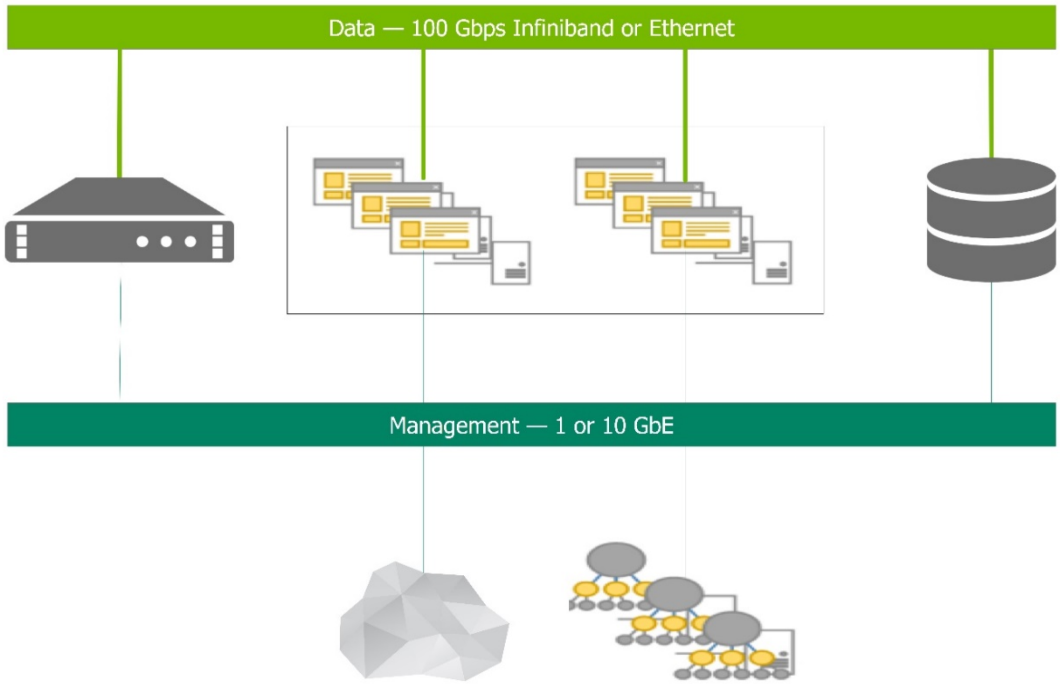

Figure 8. DGX POD networking

Whether deploying single or multiple DGX PODs, it is best to use tools that take care of the management of all the nodes. Use the NVIDIA AI software stack to provide an end-to-end solution from server OS installation to user job management.

Installing the NVIDIA AI software stack on DGX PODs requires meeting a basic set of hardware and software prerequisites, setting local configuration parameters, and deploying the cluster following step-by-step instructions. Optional installation steps include configuring multiple job schedulers and integrating with enterprise authentication and storage systems.

Once the DGX POD racks are powered-on and configured, verify all the components are working correctly. This can be an involved process. During this operational checkout, the following should be verified:

- Compute hardware, networking, and storage are operating as expected
- Power distribution works properly under maximum possible load
- Additional site tests as may be required for the data center

Day-to-day operation and ongoing maintenance of a DGX POD is greatly simplified by the NVIDIA AI software stack. Typical operations such as installing new software, performing server upgrades, and managing scheduler allocations and reservations can all be handled automatically with simple commands.

Summary

The DGX POD for Oil and Gas Predictive Maintenance reference architecture provides a blueprint to simplify deployment of GPU computing infrastructure to support the monitoring of the diverse set of assets in today's O&G operations ensuring maximum uptime through real-time anomaly detection and predictive maintenance recommendations. By leveraging today's state of the art DL algorithms accelerated by the NVIDIA AI software stack, customers gain improved insight into their asset health and operations, reduced costs associated with reduced equipment downtime and lost revenue.

A single DGX POD supporting small workgroups of AI developers can be grown into an infrastructure supporting thousands of assets. This reference architecture is based on the NVIDIA DGX SATURNV AI supercomputer which has over 1000 DGX servers and powers internal NVIDIA AI software R&D including medical imaging, autonomous vehicle, robotics, graphics, HPC, and other domains.

The DGX POD has been design and tested by using specific storage and networking partners. In addition to those mentioned in this paper, NVIDIA is working with additional storage and networking vendors who plan to publish DGX POD-compatible reference architectures using their specific products.

Finally, this white paper is meant to be a high-level overview and is not intended to be a step-by-step installation guide. Customers should work with an NVIDIA Partner Network (NPN) provider to customize a configuration and installation plan for their organization.

Legal Notices and Trademarks

Notices

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation (“NVIDIA”) does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer. NVIDIA hereby expressly objects to applying any customer general terms and conditions regarding the purchase of the NVIDIA product referenced in this specification.

NVIDIA products are not designed, authorized or warranted to be suitable for use in medical, military, aircraft, space or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on these specifications will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do the necessary testing for the application to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this specification. NVIDIA does not accept any liability related to any default, damage, costs or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this specification, or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this specification. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA. Reproduction of information in this specification is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright