



Industrializing AI for the Enterprise with NVIDIA DGX Systems and MLOps

White Paper

Document History

DGX-MLOPS v01

Version	Date	Authors	Description of Change
01	October 2021	NVIDIA	Initial release

Table of Contents

Industrializing AI with MLOps	4
The Impact of Model Debt.....	5
Taking a Page from Web Application Development	5
Managing Resources, Experiments, Jobs, and Workflow.....	6
Speeding up the End-to-End Lifecycle of AI Development.....	7
Leveraging MLOps with NVIDIA DGX-Ready Software	8
Industrialization is the Key to AI Success.....	9

Industrializing AI with MLOps

As projects leveraging artificial intelligence (AI) continue to prove their value, organizations find they have an increasing number of teams wanting to use AI to meet their business objectives. Since AI is compute intensive and working with underpowered hardware creates severe impediments to a project's success, organizations invest in accelerated computing infrastructure—including hardware and software purpose-built for AI—to shorten time-to-solution and get their models into production sooner.

Beyond good infrastructure planning, to meet the challenges of an end-to-end strategy of delivering models from prototype to production, CIOs and data science leaders need to develop cross-functional strategies that bridge the gap between traditional IT operations and AI development. These highly skilled teams share the same goal of creating successful AI-based solutions, but their pain points and daily circumstances are very different. Fortunately, a relatively new set of principles known as *MLOps* can make this process much easier, and help teams avoid many of the common pitfalls that impact unsuccessful AI initiatives.

The Impact of Model Debt

As AI-based projects continue to proliferate, a perhaps unexpected inefficiency has now become a major concern for CIOs and data science leaders: many models, despite best efforts, simply never make it to production. There are a variety of reasons for this, ranging from the mundane such as projects being cancelled, to the more serious, where the model design or available training data doesn't lead to sufficient real-world accuracy to be useful.

These abandoned, zero-value models bring down the efficiency of an AI organization, leading to what is often called *model debt*. A bit like maintaining monetary debt, model debt is a sunk cost that must be repaid someday, dragging down the profitability of the project. Organizations are beginning to understand that avoiding model debt is a key to success using AI.

A major contributor to model debt is ineffective resource and workflow management which makes it difficult to guide a model from prototype through to successful deployment. As multiple teams work on projects on shared AI infrastructure, managing access becomes challenging. Teams find themselves using shared calendars, spreadsheets, and instant messaging to resolve scheduling conflicts. As AI teams expand, day-to-day management needs such as monitoring, reporting, and job prioritization, become even more complex. This leads to teams scrambling for time on the hardware—sometimes reserving one or more systems indefinitely even when they don't have jobs actively running just to ensure they have access when they need it. In some cases, these reservations are held inactive for long periods of time—or worse, forgotten about—effectively throwing capacity into the waste bin. There is no central clearing house where decisions on the prioritization of users and jobs for the available hardware can be made.

Another related concern is the inability to easily track data and artifacts (e.g., datasets, model versions, and experiment results) for each project. This is also often handled manually, creating massive inefficiencies in larger organizations.

At its core, this is an IT problem that needs to be solved for the entire organization, much as IT is responsible for email service, HR apps, and web infrastructure, with the AI practitioners as IT's customer. Unfortunately, the infrastructure that is used for AI is often treated as bespoke, with the management burden landing squarely on the shoulders of the data science team with little or no IT and operations involvement.

Taking a Page from Web Application Development

More than a decade ago, web application development was in a similar situation. The number of web application projects was increasing rapidly but were divorced from the well-tested processes used by operations teams to manage deployment. These disconnects slowed organizations down, extending the development and deployment process.

The fix for this was something we now know as DevOps. In DevOps, the development processes used by developer teams and the operations processes used by IT were merged to create a cross-functional and repeatable set of practices to ensure rapid development and deployment of web applications.

While AI projects have different characteristics than web applications, something very valuable can be learned from DevOps. The key to shortening time to delivery is to make the development and operations processes peers, each with its own set of equally important aspects that can overlap. Identifying these overlapping areas and using the best practices of each discipline to ensure they fit

together harmoniously—rather than crashing into each other destructively—dramatically improves outcomes.

A similar situation exists for AI. Managing workflow is a common task spanning development and operations teams. As mentioned above, many AI projects are managed manually, which obviously isn't an operations best practice. However, by adopting operations techniques and adding workflow management software into the process, AI teams can increase their productivity significantly. This is *MLOps*, a way of automating and operationalizing end-to-end machine learning (ML) and AI pipelines across data preparation, model building, deployment, and production, and it plays a vital role in the success of AI projects.

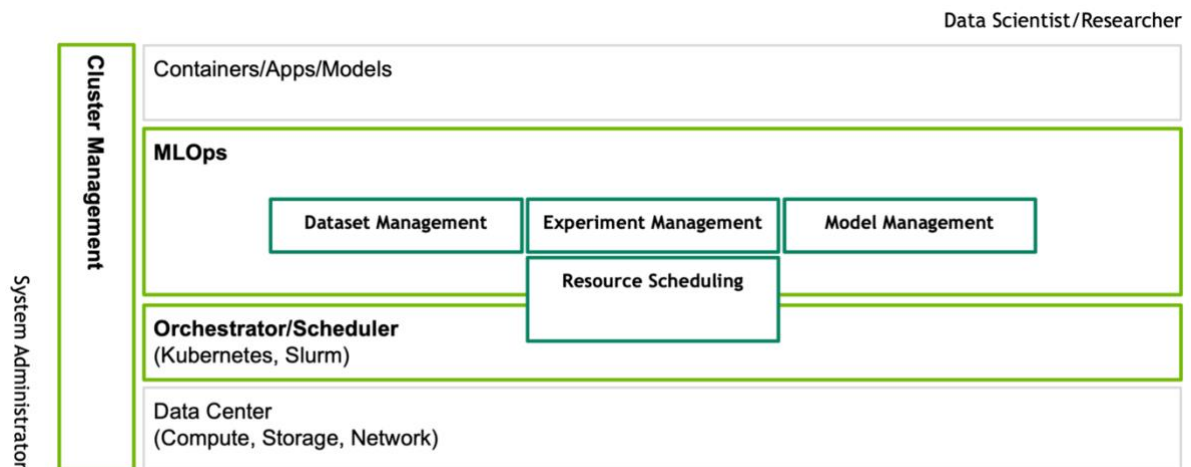
Managing Resources, Experiments, Jobs, and Workflow

MLOps tools bring an organized, methodical discipline to AI project operation and deployment, helping to ensure AI infrastructure is well utilized. They bring together the experiments, GPU resources, and specific AI jobs into a singular management environment to:

- Improve user productivity and speed workflow - users can focus on creating their models rather than handling version control or negotiating with other teams for access to resources.
- Maximize utilization of resources - keep systems busy, run more jobs than would be possible manually.
- Allow projects to scale easily - give users instant access to 1 GPU or hundreds, right-sizing every job based on priority and need.
- Accelerate the ROI of AI - ensure models get from concept to production efficiently.

MLOps software offers capabilities such as interactive sessions, dataset and experiment management, full pipeline management with model versioning, job scheduling, reporting, and much more.

Figure 1. MLOps infrastructure stack.



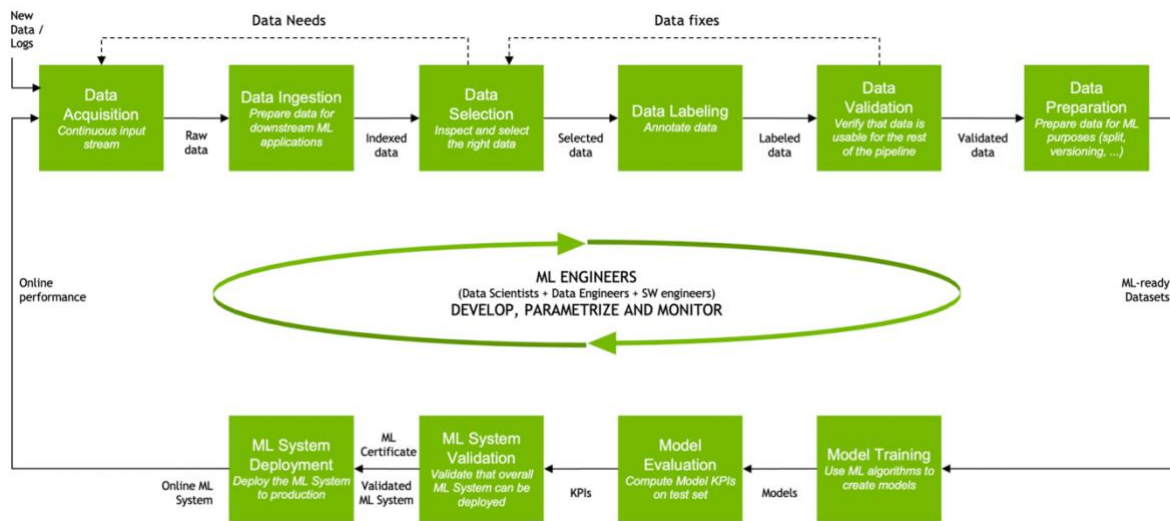
Many of these application-level MLOps tools rely on an underlying orchestration system. Two of the most common tools for this purpose are Kubernetes (K8s) and Slurm. These platforms are deployed

onto a cluster to manage the hardware resources and intelligently schedule workloads and jobs. This provides researchers a direct interface to run their experiments without needing knowledge of the underlying cluster. It also provides the MLOps engineers a tool to launch services that extend the capabilities of the base platform. For example, Kubernetes and Slurm are primarily command line driven, but MLOps tools can be deployed to give researchers a UI driven platform. There are many cloud-native products that can be easily installed onto any Kubernetes cluster and many of these are members of the NVIDIA DGX-Ready Software program discussed below.

Speeding up the End-to-End Lifecycle of AI Development

AI development is an iterative process, where data scientists continually refine their models hundreds or thousands of times to improve their accuracy, which, in turn, increases efficacy. And due to very large datasets and algorithmic complexity, the AI development process is extremely resource intensive.

Figure 2. Typical AI development cycle.



Even with a well-planned MLOps-based approach, software can't be successful in isolation. Hardware and software need to work in harmony, each as a pillar supporting the AI development platform. Organizations need a systemized, enterprise-grade approach that supports fast prototyping, frequent iteration, and continuous feedback, as well as a robust infrastructure that can scale in an enterprise production setting.

With NVIDIA DGX™ systems, featuring NVIDIA A100 Tensor Core GPUs and NVIDIA InfiniBand networking, NVIDIA has brought together over a decade of AI experience and know-how, integrated into a full-stack platform, purpose-built for the world's most complex enterprise AI challenges. The [record-breaking performance](#) of NVIDIA DGX A100 and NVIDIA DGX SuperPOD™ and the fully accelerated DGX software stack pair perfectly with MLOps solutions, ensuring that each trip through the development cycle is fast and efficient. This results in a faster time-to-solution, as you would expect. It also results in better model accuracy, since developers can use extremely large data sets

and perform more optimizations and fine tuning. This translates directly into business value, by increasing the quality of AI models and the services built around them.

Figure 3. NVIDIA DGX systems deliver record-breaking performance for AI workloads.




Choosing the right combination of MLOps software and AI infrastructure can be the difference between projects languishing in development—with developers standing idle waiting for their jobs to complete—or successfully meeting business objectives in production. NVIDIA DGX systems speed the end-to-end lifecycle of enterprise AI across development, deployment, and on-going optimization.

Leveraging MLOps with NVIDIA DGX-Ready Software

The NVIDIA DGX-Ready Software program is a portfolio of ISV solutions certified by NVIDIA that help organizations using NVIDIA DGX systems improve data science productivity, accelerate AI data pipeline and workflow, and improve accessibility and utilization of resources. The program features a variety of software solutions covering MLOps, cluster management, and resource scheduling and orchestration.

Figure 4. The ecosystem of NVIDIA DGX-Ready Software partners that can be used for the various parts of an AI infrastructure solution.



Since DGX-Ready Software is validated and certified with DGX systems, it takes the guesswork out of deploying this critical management layer for data science teams. The program features a large selection of MLOps software, allowing users to evaluate a variety of solutions knowing that all of them are designed to work in large, multi-node environments with DGX systems.

Industrialization is the Key to AI Success

Now that AI is common in day-to-day operations for many organizations and projects continue to multiply, it's clear that successfully leveraging AI requires a well thought out plan to guide projects from development through to deployment. CIOs and data science leaders know that critical parts of the equation include users with AI expertise, powerful NVIDIA GPU hardware such as NVIDIA DGX A100, fast networking such as NVIDIA InfiniBand, and optimized AI software and frameworks such as those from NVIDIA NGC.

Beyond the infrastructure, organizations are turning to MLOps software to provide powerful tools to bring IT/operations and data science teams closer together, helping them to scale AI infrastructure more easily, streamline their processes, and avoid model debt. The NVIDIA DGX-Ready Software program certifies MLOps and other software solutions to ensure they run optimally on DGX systems, allowing enterprises to leverage this new breed of AI software with confidence. This combination of NVIDIA DGX infrastructure and certified MLOps software provides a clear path for enterprises to industrialize AI, unlocking the full potential of their AI initiatives.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. 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 document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information 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 (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

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 this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order 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 document. NVIDIA accepts no 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 document 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 document. 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 document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND 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. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. 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 Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, NVIDIA DGX, NVIDIA DGX A100, NVIDIA DGX SuperPOD, NGC 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

© 2021 NVIDIA Corporation. All rights reserved.

