



# NVIDIA SITE RELIABILITY ENGINEER

Ensuring Optimal Operations and Performance for NVIDIA® DGX SuperPOD™

Running a cluster can be a complex endeavor.

A deep understanding of the hardware architecture, the software stack, and your organization's unique needs must all be considered on a daily basis. To simplify your operations, consider working with an Site Reliability Engineer (SRE), an NVIDIA DevOps engineer who works remotely to train your staff on how to manage and maintain your cluster.

As an advocate for your needs, your NVIDIA SRE:

- > **Provides guidance on cluster management.** including how to manage scheduling software, such as Slurm, and monitor the state of individual nodes. They also work with your NVIDIA Technical Account Manager (TAM) to resolve cluster operation issues.
- > **Adapts to your organization** and its unique needs for the cluster, including fine-tuning the deployed software to ensure that jobs are performing as expected and specific applications can be run.
- > **Offers expert insight** into your ML/OPs deployment and works with you to ensure the infrastructure is designed to accommodate future workloads.

## What are the Benefits of the SRE Program?

- > Ensures that the day to day cluster operations are run reliably and efficiently
- > Accelerates the transfer of knowledge with your IT staff
- > Trains your staff on how to perform SRE services for high performance, large scale DGX infrastructure based on NVIDIA best practices
- > Provides a comprehensive understanding of the tools used for monitoring and managing workload performance and cluster utilization

## Who Can Purchase SRE Services?

- > SRE services are reserved exclusively for DGX SuperPOD customers
- > A premium TAM (PTAM) contract is a prerequisite for this service contract

To learn more, contact [enterpriseservices@nvidia.com](mailto:enterpriseservices@nvidia.com)