

Platform Guidelines for GTC Content

Thanks so much for your interest in delivering hands-on training at GTC!

This document covers how to get your content ready for student interaction at GTC. Please read this document carefully, and in its entirety.

You will receive contact info for a member of the DLI Team who will be the Person in Charge (PIC) of seeing your content into production readiness and making the final approval or rejection of your submission. After reading this document, as you complete the steps described in it, or if you have additional questions, please reach out to dli-content@nvidia.com.

High Level Steps

This section outlines the high level steps required to provide students your content in an interactive infrastructure. The remainder of the document will provide more detailed steps, and will conclude with a section of frequently asked questions.

- 1) Opt in or out of using the DLI Platform. If you choose to opt out of using the DLI Platform (and will instead provide your own infrastructure for the workshop), please skip to step 5.
- 2) Prepare your content to be integrated into the DLI Platform (detailed instructions below).
- 3) The DLI Team will port your content to the DLI Platform and give you access to your content on the DLI Platform.
- 4) (optional) Continue developing content from within the provided DLI Platform environment.
- 5) Deliver a 2 hour (remote) dry-run of your content to the DLI team.

1 - Opting In or Out of the DLI Platform

The DLI Platform can provide you and your students with virtual machines on which to run your content in Jupyter (either Jupyter Notebooks or JupyterLab). We provide the option of virtual machines with one A10 or T4 GPU running CUDA 12. If you feel you need a larger GPU (e.g. A100) or multiple GPUs, please reach out to your DLI PIC to discuss.

Students will nominally be provided 1M bandwidth each on the wifi at GTC. This is enough for interfacing with Jupyter notebooks on the DLI platform, but may be inadequate for more complex graphical interfaces.

Alternatively, if you are willing, or need, to provide your own infrastructure, you can opt out of using the DLI Platform. Please note you will be required to deliver a dry-run (step 5 above)

using your provided infrastructure, which should demonstrate the ability to support interactive training for an entire class of students.

2 - Prepare Content for DLI Platform Integration

If you choose to opt in to using the DLI Platform, the DLI team needs the following from you in order to integrate your content into the DLI Platform:

- Access to a git repository (GitHub, GitLab etc.) containing all assets (except large data files, see below) required for your content.
- A Docker container using assets in the repository to run your content. The Dockerfile used to build the image for this container should reside in your repository. If your content requires the use of multiple containers, we support Docker Compose, and your repository should contain all assets associated with running the multiple containers used in your Docker Compose configuration files.
- Access to any large data files you might need for your content, provided in some method convenient for you (Google Drive, S3 bucket, etc).
- A README.md in your repository which contains:
 - The instructions to run the Docker container(s) from the image built by the Dockerfile(s) and/or docker-compose.yml in your repo.
 - (If you have any large data files) the location within the container where these data files are expected to reside.

VERY IMPORTANT: The DLI should be able, using your repository, its Dockerfile(s), and the instructions in the README, to run a containerized version of your content, accessed via a web browser, at which point your content should run successfully, and without errors, as expected. **If any of this is not true, your DLI PIC will be reaching out to you to ask you to address these issues. We expect your content to be usable as just described by the DRAFT content submission deadline. Failure to do so can result in a cancellation of your session.** Again, if you need help with this process, reach out to your DLI PIC **early**.

For a very simple example of a working repository, ready to be ported to the DLI Platform, please visit the [GTC Example Repo](#).

3 - The DLI Team Ports your Content to the DLI Platform

After you successfully deliver the above to the DLI **prior to the DRAFT content submission deadline**, the DLI team will port your content to the DLI Platform.

4 - (optional) Continue Developing from Within the DLI Platform Environment

After your content has been ported by the DLI Team to the DLI Platform, the DLI will then provide you with instructions for how to further develop your draft content (if needed) in the DLI Platform environment so that we don't need to do additional integration steps after the initial port of your content to the DLI Platform.

Once your content is ported to the DLI Platform (step 3) it is your responsibility to maintain the content in a working state.

5 - Deliver a Dry Run of your Content

Either on the DLI Platform, or on your own infrastructure (depending on whether or not you opted in to use the DLI platform), you will remotely deliver a 2 hour dry run of the content to the DLI Team. Consider this a dress rehearsal so that we can ensure your content is complete and ready to be delivered successfully at GTC. This dry run will be recorded to provide guidance on your course to TAs helping out at GTC.

After the successful completion of this step, your content can receive final approval for GTC.

FAQs

What do I do if I have any questions during this process?

Reach out to the DLI Person in Charge assigned to you.

Does my content need to run in a Jupyter notebook?

The DLI Platform expects all students to interact with your content via their web browser. For this reason we have seen almost exclusively the use of Jupyter to deliver content on the DLI Platform.

The DLI Platform supports the use of both Jupyter Notebooks and JupyterLab, so take your pick.

What do I do if I don't want to run my content in a Docker container?

Running content in a Docker container is a requirement of the DLI Platform. If you cannot or do not wish to containerize your content, you will need to provide your own infrastructure.

What do you mean by “instructions in the README for running the container?”

These instructions will take the form of `docker run <your-image>` along with any other command line flags, such as port provisions, volumes, security flags, the `--runtime nvidia` flag etc. See the [GTC Example Repo](#) for a viable example.

Will I or my students have access to the VM running my containerized content?

No. Only to the container’s exposed port(s) via web browser. If terminal access of some sort is needed, consider using Jupyter’s built-in terminal feature.

Do I have to travel to conduct the dry run?

No. If it is convenient for you, we can set up a video conference where you can present your 2 hour dry run.

What do I do if I need to make travel arrangements for GTC before I will be able to get final approval for my content?

Please refer to the GTC guidelines for travel.

How do I send you my large data files?

Any way that is convenient: Google Drive, public S3 bucket, etc.