

Dell Precision Data Science Workstation Guided Install Edition **Installation Manual** Release 3.0





Contents

Introduction
Pages 4-5

The way forward
– Configurations for ML/DL
Pages 6-7

Getting things started
– OS image and installation
Pages 8-9

Powering up your performance
– Setup and optimizations
Pages 10-11

Tool up
– Installing the NVIDIA®
Data Science Stack
Pages 12-13

Appendix
– Ubuntu Installation and Re-
installation Guide
Pages 14-16

Introduction

Albert Einstein is widely credited as saying, “The definition of insanity is doing the same thing over and over again, but expecting different results.” Converse to that, the beauty of machine learning is that it allows hardware to learn from previous computations to produce reliable, repeatable decisions and results. But of what use is that, if the process still has to rely on human guidance and intervention?

To turn potential into reality, Dell has developed Data Science workstations (DSWs), a curated set of Precision workstations pre-bundled with hardware and software, tailored for Machine Learning and Deep Learning developers.

Those in the know are fully aware that for Machine Learning to truly be successful and barriers removed, developers need access to the highest levels of computation power. Dell’s Precision workstations deliver the power to deploy and manage cognitive technology platforms, including Machine Learning (ML), Artificial Intelligence (AI) and Deep Learning (DL).

Therefore, Dell’s DSW systems are built with the latest NVIDIA® RTX™ GPUs (RTX™ A6000, 8000, 6000 5000), Intel® Xeon® CPUs, ECC memory, and advanced storage solutions, all available on the Dell website.

Backed by Dell’s turnkey approach to IT, and allowing developers to hit the ground running, these solutions are certified and listed as supported “NVIDIA® NGC-Ready” platforms on the NVIDIA® site.

NGC-Ready system validation includes tests of:

- Single and multi-GPU Deep Learning training using TensorFlow, PyTorch and NVIDIA® DeepStream Transfer Learning Toolkit
- High volume, low latency inference using NVIDIA® TensorRT, TensorRT Inference Server, and DeepStream
- Data Science ML using RAPIDS and XGBoost
- Application development using the CUDA Toolkit

In fact, these NVIDIA GPU-accelerated DSWs are designed to get the developer working on AI/ML/DL datasets the same day the DSW arrives from the Dell factory. With that on the table, the power is in the developer’s hands to experiment, prepare data, and train AI/ML/DL models orders of magnitude faster than possible on CPU-only AI development platforms. It couldn’t be simpler.

Theory is one thing, practice is another, so Dell has created this step-by-step installation manual, designed to help customers easily install the “NVIDIA® Data Science Software, powered by RAPIDS” onto the DSW received from the Dell factory pre-loaded with a custom Ubuntu 20.04 LTS image. A manual that will also prove incredibly useful in instructing users on how to install the NVIDIA® Data Science Software onto a user-installed Ubuntu 20.04 LTS OEM image provided by Dell.

The intent of this short guide is to help users quickly set up the OS and the NVIDIA® Data Science Software specifically on Dell DSWs. It is also intended to help them overcome the three time-consuming hurdles to getting started as a data science developer. These are:

1. Configuring a system with the right hardware components.
2. Installing an OS that assures support for all devices out-of-the box.
3. Compiling a curated software stack with all the necessary NVIDIA GPU-accelerated libraries and machine learning and deep learning frameworks.

For best practices, and step-by-step guidance on installing Ubuntu OS on your Data Science Workstation, [visit our appendix here.](#)

To get developers started using their Data Science Workstation as quickly as possible, we offer NVIDIA® NGC-Ready Dell Precision workstation configurations validated for the NVIDIA® Data Science Software.



1. The way forward – Configurations for ML/DL

Power is nothing without control, or direction. As such, Table 1 below includes recommended configurations of mobile, tower and rack Data Science workstations that are suitable for ML/DL.

The list ranges from mobile workstations and mid-sized towers to larger tower and rack workstations with with multiple CPUs & NVIDIA GPUs.

These configurations are suitable for data preparation, all types and sizes of ML/DL model training, and running the resulting inference engines. But that's not the end. [Find more configuration information by visiting the Dell Precision Workstation Solutions webpage.](#)



Precisely developed for your needs



Precision 7550:

Dell's most powerful 15" mobile workstation is AI ready, featuring the NVIDIA® Quadro RTX™ 5000 graphic card.



Precision 7750:

The most powerful Precision mobile workstation is AI ready, featuring the NVIDIA® Quadro RTX™ 5000 graphic card.



Precision 5820 Tower:

Ideal for cognitive solution development and inference applications. Features NVIDIA® RTX™ A6000 or 8000 graphic cards.

Precision 7920 Tower and Rack:

Handles learning model training and larger solution frameworks with ease. Features NVIDIA® RTX™ A6000 or 8000 graphic cards. And give up no performance when you remote log into the 7920 Rack from home.



2. Getting things started – OS image and installation

This document aims to provide robust installation guidance, even as software and script versions evolve. However, steps may change in between guide releases. Always check [here](#) for late-breaking installation updates.

Need to reinstall your Data Science Workstation's operating system? First you will need to confirm that "Secure Boot" is disabled in the system BIOS/UEFI before proceeding. The NVIDIA® Data Science Software installation requires Secure Boot to be disabled to support LINUX kernel updates.

Solve practical machine learning and deep learning problems with any one of these Dell Precision workstations with extensive memory, outstanding CPU and with multiple CPUs & NVIDIA GPU processors, and performance storage to power advanced cognitive solutions.

	NVIDIA® RTX™ A6000	NVIDIA® Quadro RTX™ 8000	NVIDIA® Quadro RTX™ 6000	NVIDIA® Quadro RTX™ 5000	Intel® Xeon® CPU	System Mem	Drives (1TB-3TB)
Precision 7920 Tower	Single/ Dual/ Triple	Single/ Dual/ Triple	Single/ Dual/ Triple	(n/a)	Dual Xeon® Platinum/ Gold	192GB- 384GB ECC RDIMM	NVMe PCIe Class 50 SSD
Precision 7920 Rack	Single/ Dual/ Triple	Single/ Dual/ Triple	Single/ Dual/ Triple	(n/a)	Dual Xeon® Platinum/ Gold	192GB- 384GB ECC RDIMM	NVMe PCIe Class 50 SSD
Precision 5820 Tower	Single/ Dual	Single/ Dual	Single/ Dual	(n/a)	Single W-2245/ W-2255/ W-2275	64GB- 256GB ECC RDIMM	NVMe PCIe Class 50 SSD
Precision 7750	(n/a)	(n/a)	(n/a)	Single	Single E-2286	64GB- 128GB ECC RDIMM	NVMe PCIe Class 50 SSD
Precision 7550	(n/a)	(n/a)	(n/a)	Single	Single E-2276	64GB- 128GB ECC RDIMM	NVMe PCIe Class 50 SSD

Dell Precision DSW systems ship with Ubuntu 20.04 LTS preinstalled, so already have an OS image pre-configured to support the hardware out-of-the-box and "Secure Boot" disabled by default. You can skip these steps and go directly to Section 3.

If your system came from the Dell factory with Ubuntu OS pre-installed and for some reason you need to recover and re-install that Ubuntu OS, [follow these instructions](#), then proceed to Section 3.

If you choose to install a different Ubuntu OS image on your system, [simply follow the best practices instructions provided in the Appendix "Installing Ubuntu OS"](#), then proceed to Section 3.



3. Powering up your performance – Setup and Optimizations (optional)

To ensure that you are extracting every last drop of power out of your Dell Precision workstation, we recommend the following optimizations to increase the performance and usability of the system for Machine Learning and Deep Learning workloads.

3.1 Disable idle user power management configuration

To complete this action, we recommend disabling power management settings for Idle User case, since in most cases during installation and development, the machine will download the software/dataset over the internet or run ML/DL workflow unattended for an extended period of time.

Disable power management settings with the command below.

```
/usr/bin/gsettings set org.gnome.settings-daemon.plugins.power sleep-inactive-ac-timeout 0
```

3.2 Disable Hibernation settings

There is no advantage in enabling “hibernate” on Dell Data Science workstations, quite simply because it is unnecessary and consumes valuable storage space. Disable it as per below.

```
sudo systemctl mask hibernate.target
```

3.3 Remove swap file

In addition, we recommend removing the “swap file”. Swap file sizes are typically 1.5x the size of installed RAM. On large RAM systems this can be a prohibitive amount of SSD storage

3.4 Is swap active?

```
sudo swapon --show
```

If it is, it will show something similar to the below

```
NAME      TYPE  SIZE USED PRIO
/swapfile file 215.3G  0B  -2
```

3.5 Deactivate the swap

```
sudo swapoff -v /swapfile
```

3.6 Remove it from /etc/fstab

```
sudo sed -i '/^\s*\/swapfile/d' /etc/fstab
```

3.7 Delete the actual file

```
sudo rm /swapfile
```



4. Tool up – Installing the NVIDIA® Data Science Stack

The NVIDIA® Data Science Software Stack is a set of tools that makes it easier to set up your Dell Data Science Workstation and manage NVIDIA®'s curated software stack for NVIDIA GPU accelerated Data Science workloads.

The NVIDIA® Data Science Stack has been tested by Dell.

4.0 The download and configuration of the NVIDIA® Data Science Stack on the Dell Data Science Workstation.

Note: The NVIDIA® Data Science Software installer will update your NVIDIA GPU driver as required.

DOWNLOAD THE NVIDIA® DATA SCIENCE STACK PACKAGE

Please refer to the following [link](#) for release notes disclosing changes and enhancement details for the different releases.

Note: Unless instructed otherwise do NOT use 'sudo'.

4.0 (a) Install git on your system using the following two commands.

1. `sudo apt update`
2. `sudo apt install git`

4.0 (b) Create the DellDSW subdirectory, download and unpack the NVIDIA® Data Science Stack in the new DellDSW subdirectory using the following commands.

1. `mkdir ~/DellDSW ; cd ~/DellDSW`
2. `git clone https://github.com/NVIDIA/data-science-stack.git`

4.0 (c) Begin the configuration of the stack on the system. Remain in the DellDSW subdirectory and execute the following commands. The final step will reboot the system to apply the changes.

1. `cd data-science-stack`
2. `./data-science-stack setup-system`
 - user may be asked to select appropriate keyboard version (choose based upon your region).
 - Also, if asked about modified configuration file grub, select "keep the local version currently installed".
3. `sudo reboot 0`

4.0 (d) Upon rebooting, return to the subdirectory where you extracted the data science stack and execute the following commands to add a user to the Docker Group.

1. `cd ~/DellDSW/data-science-stack`
2. `./data-science-stack setup-user`

4.0 (e) Logout and log back into the system using the following command.

1. `gnome-session-quit --no-prompt`

4.0 (f) Once you have logged back in, change to the directory where the data science stack was extracted.

1. `cd ~/DellDSW/data-science-stack`

4.0 (g) Proceed to the next section of this guide.

4.1 Now either choose to create a container with the data science development environment or install it natively on your system. Alternatively, do both.

Based on your choice, follow the instructions below:

4.1 (a) Option 1: Build a container:

1. `./data-science-stack build-container`

4.1 (b) Option 2: Create a Conda environment for NVIDIA® data science distro:

1. `./data-science-stack build-conda-env`

4.1 (c) Logout and log back into the system using the following command.

1. `gnome-session-quit --no-prompt`

4.1 (d) Change to the directory where the data science stack was extracted.

1. `cd ~/DellDSW/data-science-stack`

Continued →

4. Tool up – Installing the NVIDIA® Data Science Stack

The NVIDIA® Data Science Software Stack is a set of tools that makes it easier to set up your Dell Data Science Workstation and manage NVIDIA®'s curated software stack for NVIDIA GPU accelerated Data Science workloads.

4.2 Once installation is complete, perform the following tests to verify success:

4.2 (a) To verify containerized environment, run:

```
data-science-stack run-container
```

If the scripts run successfully creating the environment, then the installation was successful.

4.2 (b) To verify conda environment, run:

```
data-science-stack run-jupyter
```

If the scripts run successfully, then the installation was successful.

And that is it. Your Dell Precision Data Science Workstation is ready for use. It's now time to move your business forward by being able to power amazing outcomes.



Appendix: Ubuntu Installation Guide

This appendix contains the simple steps and best practice guidance for installing an Ubuntu OS onto a bootable storage drive on a Dell Data Science Workstation. It is intended for users wanting to do their own installation of Ubuntu instead of using the Ubuntu OS pre-installed at the Dell factory. Use and adapt these steps as a best practice to install a different Ubuntu OS version of your own choosing, although it would not be an installation already validated and optimized by Dell, NVIDIA®, and Canonical. Dell recommends using Long Term Support (LTS) Ubuntu versions only.

If the workstation came from the Dell factory with Ubuntu pre-installed, **please refer to these instructions** to recover the factory installed operating system.

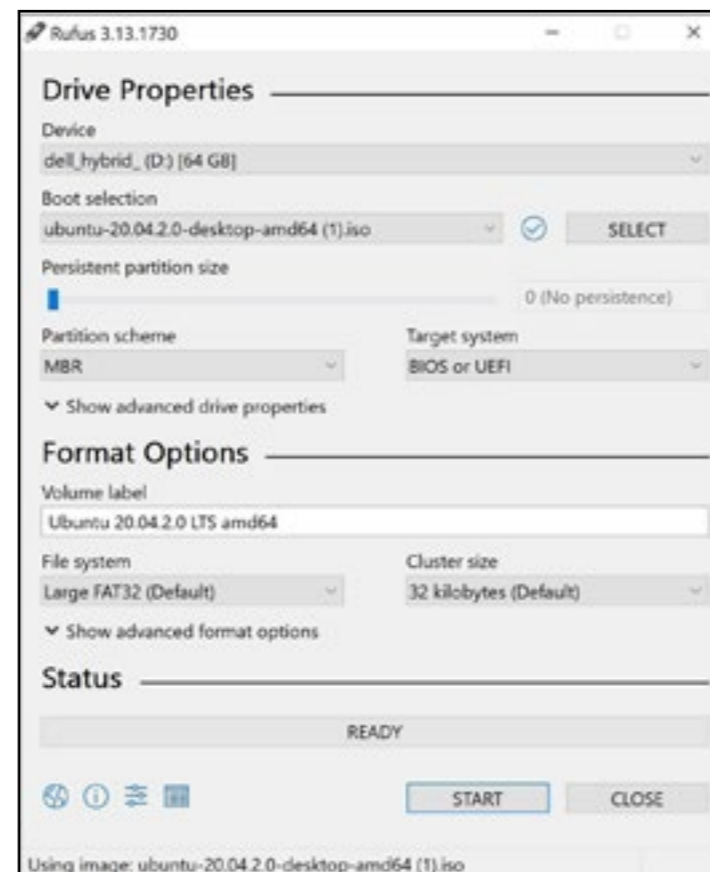
These steps require reformatting the storage drive and will wipe out any existing drive content, so be sure to save any important drive content before starting.

Unless instructed otherwise, do NOT use 'sudo', and use a standard user account in a Bash Linux shell (i.e., terminal window).

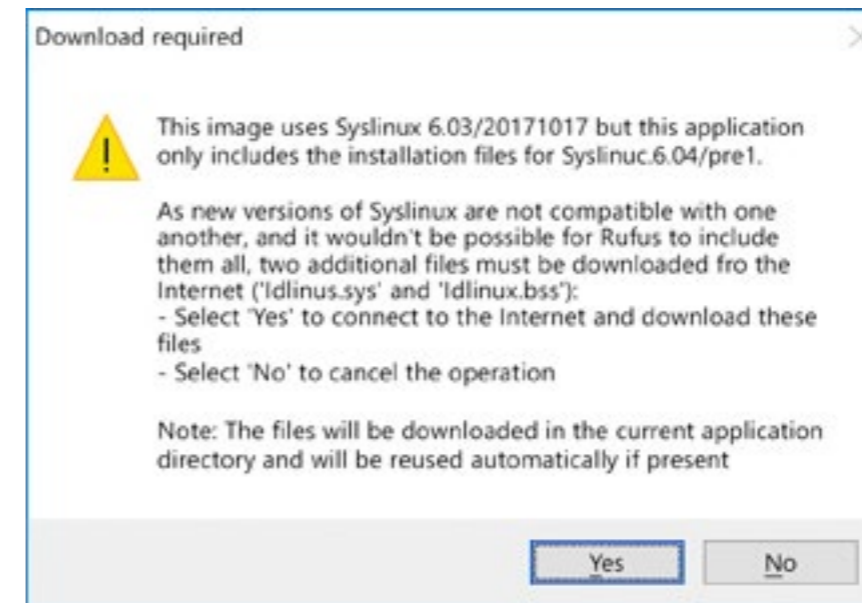
1.0 Download Ubuntu 20.04 media [here](#).

2.0 (a) Create installation media

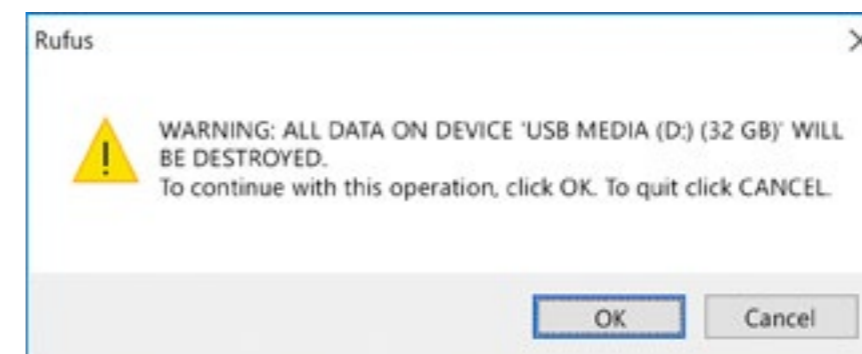
1. Download and run [Rufus](#).
2. Insert your USB media device.
3. Configure as shown below and click **Start**.



2.0 (b) Click Yes.



2.0 (c) Click OK.



Appendix: Ubuntu Installation Guide

3.0 Configure system BIOS.

- 3.0 (a) Power up or reboot system. When the Dell logo appears, immediately tap **F2** repeatedly to load the one-time boot menu.
- 3.0 (b) Disable Secure Boot
 1. Settings → Secure Boot → Select Disabled
- 3.0 (c) Change SATA setting to AHCI
 1. Settings → System Configuration → SATA Operation → Select AHCI
- 3.0 (d) If yours is a Precision Mobile Workstation, skip this step and go to step 3.0 (e).
 1. Settings → Advanced Boot Options → Check Enable Legacy Option ROMs
 2. Settings → Boot Sequence → Select Legacy
 3. *Note: This will cause systems equipped with NVMe hard disk carriages to produce a flashing yellow disk health light. This is normal and is the result of using the Legacy boot option.*
 4. *Note: If this is a Precision Mobile Workstation please keep UEFI selected.*
- 3.0 (e) Click **Apply, OK, Exit**.

4.0 Install Ubuntu 20.04.

- 4.0 (a) Insert USB installation media device created in Appendix Step 2.
- 4.0 (b) Power up or reboot system. When the Dell logo appears, immediately tap **F12** repeatedly to load the one-time boot menu.
- 4.0 (c) Under UEFI Boot, select your USB installation media device and press <Enter>.
- 4.0 (d) After the installer loads, select preferred language and click **Install Ubuntu**.
- 4.0 (e) If you boot into Ubuntu, you will need to double click on the Install Ubuntu 20.04.4 icon.
- 4.0 (f) Select keyboard layout and click **Continue**.
- 4.0 (g) If asked about WIFI, select **I don't want to connect to a WIFI network right now**.
- 4.0 (h) If asked about updates and other software, select **Minimal Installation**.
- 4.0 (i) Uncheck **"Download updates while installing Ubuntu."**
- 4.0 (j) Uncheck **"Install third-party software for graphics and Wi-Fi hardware and additional media formats."**

Note: The compatible NVIDIA® driver gets installed during the later step when the NVIDIA® Data Science stack is installed.

- 4.0 (k) Click **Continue**.
- 4.0 (l) Select **"Erase disk and install Ubuntu."**
- 4.0 (m) Click **Install Now**.
- 4.0 (n) Confirm writing changes to disk by clicking **Continue**.
- 4.0 (o) Select **time zone** and click **Continue**.
- 4.0 (p) Enter name, computer name, username, password and click **Continue**.
- 4.0 (q) Ubuntu will install for the next several minutes. When it finishes, click **Restart Now**.
- 4.0 (r) Remove USB installation media device.
- 4.0 (s) Log into the system with username and password entered during installation.

5.0 Configure Dell repositories and update all packages.

5.0 (a) Open a Terminal Window and enter the following commands:

1. `sudo apt-key adv --recv-keys --keyserver hkp://keyserver.ubuntu.com:80 F9FDA6BED73CDC22`
2. `sudo apt update`
3. `sudo apt upgrade -y`
4. `sudo add-apt-repository 'deb http://dell.archive.canonical.com/updates/ focal-dell public'`
5. `sudo add-apt-repository 'deb http://dell.archive.canonical.com/updates/ focal-oem public'`
6. `sudo add-apt-repository 'deb http://dell.archive.canonical.com/updates/ focal-somerville public'`
7. `sudo add-apt-repository 'deb http://dell.archive.canonical.com/updates/ focal-somerville-melisa public'`
8. `sudo apt update`
9. `sudo apt upgrade -y`
10. `reboot`

5.0 (b) The system will restart.

6.0 Proceed with instructions in Dell Data Science Workstation Installation Guide, Section 3 to install the NVIDIA® Data Science Software stack.



[Learn more](#) about Dell Precision workstation solutions



[Contact](#) a Dell Expert

Connect on social



Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. This installation guide is for informational purposes only.

The contents in this installation guide were accurate at the point of publication, April 2021.

NVIDIA and NVIDIA Quadro are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and/or other countries. Products may differ from images displayed.

