

QUICK START GUIDE

NI Real-Time Hypervisor

Version 2.0

The NI Real-Time Hypervisor provides a platform you can use to develop and run LabVIEW and LabVIEW Real-Time applications simultaneously on a single controller.

(Windows) Refer to the **NI Real-Time Hypervisor for Windows** section below for a summary of your system configuration.

(Linux) Refer to the **Installing the Hypervisor for Linux** section below for installation instructions.

NI Real-Time Hypervisor for Windows

Your NI Real-Time Hypervisor ships pre-configured with everything you need to develop and run LabVIEW and LabVIEW Real-Time applications simultaneously on a single controller. NI Factory Installation Services (FIS) has equipped your system with hardware that meets the requirements of the NI Real-Time Hypervisor. In addition, FIS has partitioned this hardware between Windows and the LabVIEW Real-Time OS (NI ETS) according to the preferences you specified during the ordering process and the resource partitioning requirements of the NI Real-Time Hypervisor.

However, you can assign the resources in your system to different OSes. Refer to the [Getting Started with Your NI Real-Time Hypervisor](#) section for instructions to complete this task.

Installing the NI Real-Time Hypervisor for Linux

You need the following items to install the NI Real-Time Hypervisor for Linux:

- Red Hat Enterprise Linux installation CD



Note The NI Real-Time Hypervisor supports only Red Hat Enterprise Linux.

- USB CD-ROM drive
- USB memory drive

Complete the following steps to install the NI Real-Time Hypervisor for Linux.

1. Install LabVIEW RT and VISA Server on the system.
 - a. Power on or restart the system.
 - b. Press the <Delete> key to access the BIOS setup utility on your system.
 - c. Change the boot option so the system boots into **LabVIEW RT**.
 - d. Save the changes and restart the system.
 - e. Launch Measurement and Automation Explorer (MAX) on a development computer.
 - f. Expand the **Remote Systems** tree and locate the system you plan to install the NI Real-Time Hypervisor on.
 - g. Expand the tree next to the NI Real-Time Hypervisor system.
 - h. Click the **Software** tree item and install **LabVIEW RT** and **VISA Server**.
2. Obtain identification information for your controller and chassis.
 - a. Expand the **Devices and Interfaces** tree.
 - b. Right-click your controller and select its controller family in the **Identify As** menu.
 - c. Expand the tree next to your controller.
 - d. Right-click your PXI chassis and select its model in the **Identify As** menu.
 - e. Connect to the system with an FTP client.
 - f. Navigate to the **ni-rt/system** directory.
 - g. Copy the **pxisys.ini** and **pxiesys.ini** files to a USB drive. These files identify your controller and chassis. Later, you will copy these files to the NI Real-Time Hypervisor system.

3. Start the Red Hat Enterprise Linux installation process.
 - a. Connect a USB CD-ROM drive to the system you plan to install the NI Real-Time Hypervisor on.
 - b. Place the Red Hat Enterprise Linux installation CD in the drive.
 - c. Restart the system.
 - d. Press the <Delete> key to access the BIOS setup utility on your system.
 - e. Change the boot option so the system boots into **Windows/Other OS**.
 - f. Change the boot priority so the system boots from the CD-ROM drive.
 - g. Save the changes and restart the system.
 - h. Begin the Red Hat Enterprise Linux installation process.
4. When you reach the **Partitioning Type** menu, create two Linux partitions. Leave enough free space for two more partitions, a partition at least 1 GB in size for the LabVIEW Real-Time OS (NI ETS) and a 1 GB partition for boot information, that you will create later.



Note The free space must not be in a Logical Volume Manager (LVM).

- a. Select **Create Custom Layout** from the hard drive menu.
 - b. Select **New** from the partition menu to create two new partitions. For example, you could create an ext3 partition mounted at / and a swap partition.
5. After you complete the installation and the system restarts, log in with root privileges.
 6. Connect the USB drive that contains the **paxisys.ini** and **pxiesys.ini** files to the system to identify your controller and chassis.
 7. Copy the **paxisys.ini** and **pxiesys.ini** files to the **etc** directory.

8. Use the `fdisk` utility that comes with Red Hat Enterprise Linux to partition the hard drive.
 - a. In the command prompt, type `fdisk` followed by the name of your hard drive and press the <Enter> key. For example, if your hard drive is named `/dev/sda`, type `fdisk /dev/sda`.
 - b. Press the <N> key and then the <Enter> key to create a new partition.
 - c. Press the <P> key and then the <Enter> key to designate the partition as a primary partition.
 - d. Specify the number of an unused partition and press the <Enter> key to begin creating the NI ETS partition. For example, if you only created two partitions during the Linux installation process, press the <3> key.
 - e. Start the partition at the default cylinder.
 - f. Enter a size for the partition. For example, to designate 30 GB for NI ETS, type `+30G`.
 - g. Repeat steps b–g to create a 1 GB partition for boot information.
9. Change the file systems of the partitions.
 - a. Press the <T> key and then the <Enter> key.
 - b. Enter the partition number of the NI ETS partition and press the <Enter> key. For example, if you designated partition 3 as the NI ETS partition, press the <3> key.
 - c. Press the <C> key and then the <Enter> key to format the NI ETS partition with the W95 FAT32 (LBA) file system.
 - d. Enter the partition number of the boot partition and then press the <Enter> key.
 - e. Type `21` and then press the <Enter> key to identify the partition as the boot partition.



Note `fdisk` returns an error that says `21` is an unknown partition type. This error is expected and does not indicate a problem.

- f. Press the <W> key and then the <Enter> key to save these changes.



Note Linux displays an error that says it failed to re-read the partition table but will use the new table when you reboot the system. This error is expected and does not indicate a problem.

10. Restart the system.

11. Use the `mkdosfs` utility that comes with Red Hat Enterprise Linux to name the NI ETS partition **LABVIEW_RT**.
 - a. Right-click the desktop and select **Open Terminal** from the shortcut menu.
 - b. Type `mkdosfs -n LABVIEW_RT` followed by the name of the hard drive and the partition number of the NI ETS partition and then press the <Enter> key. For example, if the NI ETS partition is the third partition of a hard drive named `/dev/sda`, type `mkdosfs -n LABVIEW_RT /dev/sda3`.
12. Place the NI Real-Time Hypervisor installation CD in the USB CD-ROM drive and run the NI Real-Time Hypervisor installer.
13. Right-click the desktop and select **Open Terminal**. Then, launch the NI Real-Time Hypervisor Manager by navigating to `usr/local/natinst/hypervisor/ConfigurationUtility` and running the `NIHypervisorManager` application. A dialog box appears and prompts you to run the setup.
14. Click the **Run Setup** button. The NI Real-Time Hypervisor manager runs the setup and prompts you to restart the system.
15. Restart the system and select **Linux** from the boot menu.
16. Log in with root privileges.
17. Compile the virtualization drivers into the Linux kernel. Refer to `Virtualization_Drivers.pdf`, located in the `usr/local/natinst/hypervisor` directory, for instructions to perform this task. The `linux-driv-src.tgz` file, which contains the virtualization drivers, is located in the `usr/local/natinst/hypervisor/system/drivers` directory.
18. Navigate to the `boot/grub` directory and open the `menu.lst` file.
19. Change the **initrd** entry to the initial RAM disk (INITRD) you created in step 18. For example, if you created an INITRD named `initrd-2.6.18-53.el5-VLX.img`, enter `initrd /boot/initrd-2.6.18-53.el5-VLX.img` for the **initrd** entry.
20. Save your changes to `menu.lst` and close the file.
21. Navigate to the `usr/local/natinst/nihypshm` directory and run `KernelDrv_Install.sh`.
22. Refer to the **Getting Started with Your NI Real-Time Hypervisor** section below for instructions to partition your devices according to your application needs and begin developing applications.

Getting Started with Your NI Real-Time Hypervisor

Complete the following steps to become familiar with your system's configuration utility, known as the NI Real-Time Hypervisor Manager, and begin developing applications.

1. Power on or restart your system.
2. Depending on your system, select **Windows** or **Linux** from the boot menu to boot into your general-purpose OS.
3. Use the NI Real-Time Hypervisor Manager to view the devices assigned to each OS.
 - a. **(Windows)** Select **Start»National Instruments»NI Real-Time Hypervisor»NI Real-Time Hypervisor Manager** to launch the NI Real-Time Hypervisor Manager.
(Linux) Right-click the desktop and select **Open Terminal**. Launch the NI Real-Time Hypervisor Manager by navigating to `usr/local/natinst/hypervisor/ConfigurationUtility` and running the `NIHypervisorManager` application.
 - b. Review the OS partition settings.
 - The **Basic** tab displays each device and the OS to which each device is assigned.
 - The **Advanced** tab displays each device, the OS to which each device is assigned, and the PCI Interrupt Request Line (PIRQ) that each device uses.
4. (Optional) Change the OS partition settings.
 - a. Click the **Basic** tab to display the list of devices and the OS to which each one is assigned.
 - b. Use the pull-down menus next to each device to change the OS to which each device is assigned.
 - c. Right-click the **Memory** row and select **Set Memory Allocation** to partition the system memory between the OSes according to your application needs.
 - d. Use the pull-down menus next to each CPU core to assign CPU cores to OSes according to your application needs.
 - e. Click the Apply button to apply the partitioning configuration. The NI Real-Time Hypervisor Manager generates a set of instructions.
 - f. Follow the instructions generated by the NI Real-Time Hypervisor Manager to ensure that each hardware device is connected to the correct card slot.
5. Power on or restart your system.
6. Select **NI Real-Time Hypervisor** from the boot menu to boot into hypervisor mode.
7. Develop and run applications just as you would with a normal host and Real-Time system.

Communicating Between Operating Systems

The NI Real-Time Hypervisor includes a virtual RT console, a virtual Ethernet connection, and a shared memory feature that you can use to communicate between the general-purpose OS and the LabVIEW Real-Time OS (NI ETS).

Using the Virtual RT Console

Use the virtual RT console to view configuration and troubleshooting information for the NI Real-Time Hypervisor. For example, you can use the virtual RT console to determine the RT target IP address. Refer to the **Connecting to the Virtual RT Console** topic on the **Contents** tab in the *NI Real-Time Hypervisor Help* for more information.

Using the Virtual Ethernet Connection

Use the virtual Ethernet connection to communicate between your general-purpose OS and NI ETS with remote communication methods. Refer to the **Transferring Data Between Operating Systems** topic on the **Contents** tab in the *NI Real-Time Hypervisor Help* for more information.

Using Shared Memory

You can share up to 95 MB of system memory between the OSes on your NI Real-Time Hypervisor system with shared memory. Use shared memory to transfer data between OSes if remote communication methods are inefficient. Refer to the **Transferring Data Between Operating Systems** topic on the **Contents** tab in the *NI Real-Time Hypervisor Help* for more information.

Restoring Your System

(Windows) Refer to the **Restoring Your NI Real-Time Hypervisor System** topic on the **Contents** tab in the *NI Real-Time Hypervisor Help* for instructions to restore your system.

(Linux) Refer to the **Installing the Hypervisor for Linux** section above to reinstall the NI Real-Time Hypervisor on your system.

Where to Go from Here

Refer to the *NI Real-Time Hypervisor Help* for complete documentation of the NI Real-Time Hypervisor, including an introduction to virtualization and techniques for communicating between your general-purpose OS and the RT target.

Refer to ni.com/info and enter the Info Code NIRTH2KI to access known issues for the NI Real-Time Hypervisor 2.0.

Refer to ni.com/info and enter the Info Code HV_Devices for information about unsupported hardware devices and drivers.

LabVIEW, National Instruments, NI, ni.com, the National Instruments corporate logo, and the Eagle logo are trademarks of National Instruments Corporation. Refer to the *Trademark Information* at ni.com/trademarks for other National Instruments trademarks. Other product and company names mentioned herein are trademarks or trade names of their respective companies. For patents covering National Instruments products/technology, refer to the appropriate location: **Help»Patents** in your software, the `patents.txt` file on your media, or the *National Instruments Patent Notice* at ni.com/patents.