
NI-VXI Software Enhancements for the VXI/VME-PCI8000 Kit for Microsoft Operating Systems

Thank you for purchasing the NI-VXI bus interface software from National Instruments. These release notes describe new features of the NI-VXI bus interface software along with enhancements of existing features.

New Features and Terminology

New functionality has been added to NI-VXI in four major areas to exploit features in MXI-2 and the MITE ASIC. These features are as follows:

- Window mapping
- DMA
- Shared memory
- Remote controllers

Window Mapping

The MITE architecture allows much more flexibility in low-level mapping of VXI address spaces. In particular, the CPU-MXI interface of the MITE has windows that can be dynamically resized and relocated from CPU space to VXI space. The NI-VXI low-level functions have new extensions that reflect this feature. Refer to Chapter 6, *Low-Level VXIbus Access Functions*, in the *NI-VXI Software Reference Manual for C* for more information.

As in earlier versions of NI-VXI, `MapVXIAddress()` checks whether a window that can be shared already maps to the desired address space and location. If so, it returns a pointer to that window. If the desired space is not already mapped, `MapVXIAddress()` sets up a new MITE window to the VXI address and returns a pointer to the new window.

The success of this allocation depends on the availability of three factors:

- MITE windows
- CPU address space
- Memory for allocating data structures for the map

MITE Windows

The MITE has eight CPU windows. NI-VXI uses three of these windows, leaving five for user applications.

CPU Address Space

The PCI-MXI-2 can decode any 32-bit address on the PCI bus as a VXI cycle, giving 4 GB of addressability, which can be used for windows on the PCI-MXI-2. The operating system or computer architecture may limit which addresses can be assigned to the PCI-MXI-2.

- ◆ **DOS users**—Under DOS, the windows must be located below the 1 MB boundary because DOS can access only 1 MB of memory. The user window under DOS can be only as large as the memory available in the upper memory block, usually not more than 128 KB.

Memory for Allocating Data Structures

Memory is needed to set up the necessary page tables. If you request a very large window—hundreds of megabytes, for example—you may run out of memory.

- ◆ **Windows 3.1 users**—A large number of segment descriptors may be needed to map the CPU address space.

The `MapVXIAddressSize()` function is the standard mechanism for specifying how large a window the driver should map on a call to `MapVXIAddress()`. The default size of a mapped window is 64 KB. The `MapVXIAddressSize()` function is described later in this document.

DMA

The MITE has two DMA channels, which NI-VXI uses to improve the throughput of block transfers to and from the VXI system. The DMA channels can use various high-speed bus protocols, such as the following:

- MXI block
- MXI synchronous
- Burst mode (on the PCI bus)
- VME64 (on the VXI bus)

The DMA channels can transfer data between a VXI device and local memory, or between VXI devices. The DMA channel can handle contiguous or noncontiguous local memory. If it is handling noncontiguous memory, it can perform scatter-gather operations on the noncontiguous memory.

The `VXImove()` function automatically uses appropriate bus protocols and transfer types to efficiently perform the data transfer specified in the function. You can also use some extra configuration options and function parameters to instruct the NI-VXI software to use DMA channels for particular types of operations and to designate what protocols the channel should use. See Chapter 7, *High-Level VXIbus Access Functions*, in the *NI-VXI Software Reference Manual for C* for complete descriptions of `VXImove()` and other high-level functions. However, previously written NI-VXI code will use the DMA capabilities without modification.

To take full advantage of the throughput of the DMA channels, you should perform 32-bit transfers where both the source and the destination are longword aligned. If you need to transfer character data between devices of different byte orders—for example, between a big-endian device and an Intel 80x86-based Windows NT PC—transfer the data as longwords but adjust the byte-ordering parameters in `VXImove()` to get the correct data in the most efficient manner.

Examples:

```
/* Transferring 32-bit data to a big-endian A32 device */
VXImove(0x0, userBuffer, 0x3, deviceOffset, numDataPoints, 4);

/* Transferring 8-bit data to a big-endian A32 device */
VXImove(0x80, userBuffer, 0x3, deviceOffset, numDataPoints / 4, 4);
```

Shared Memory

You can share a portion of your system RAM in VXI space, as on previous drivers. In addition, MITE-based boards have SIMM slots for installing onboard RAM. On a PCI-MXI-2, this RAM can be shared in a VXI address space. Furthermore, you can divide the VXI space assigned to your device into two halves, sharing the onboard RAM in one half and system RAM in the other. The configuration options, such as byte-ordering, can be different for each half of the region of VXI space you are sharing. For more information review the options of the PCI-MXI-2

Logical Address Configuration Editor in Chapter 6, *NI-VXI Configuration Utility*, of *Getting Started with Your VXI/VME-PCI8000 and the NI-VXI Software for Microsoft Operating Systems*.

Remote Controllers

If you are using a VXI-MXI-2 VXIbus extender in your system, you have remote controller capability. Each VXI-MXI-2 has a full MITE chip set onboard. Although it does not have a CPU connected to it, this MITE chip set gives the VXI-MXI-2 many of the same features as the PCI-MXI-2, including onboard memory, DMA channels, TIC-based triggering, and interrupt mapping. This means that a MXI-2 system has at least one full-featured bus controller residing in each VXI chassis, controlled by and communicating with the local controller via the VXIbus and the MXIbus. These MITE-based bus controllers are called *remote controllers*.

In previous versions of NI-VXI, the CPU-MXI-based controller was called an *external controller*, and VXI-MXI boards on the same MXIbus as the external controller were called *extending controllers*. The extending controllers and the external controller together made up one *extended controller*. The presence of the full-featured remote controllers has led to a slight change in this model. The CPU-MXI is still an external controller, but instead of having extending controllers, each parent-side VXI-MXI-2 is a separate remote controller. Other VXI-MXI-2 boards are treated simply as extenders that may have additional memory present, but may not perform bus access, trigger, or interrupt services. This definition assigns one remote controller to each chassis, resolving any resource arbitration issues for the system.

Any NI-VXI function that previously accepted a **controller** parameter—whether extended, external, embedded, or local—will now accept a remote controller as well. A new option in VXIedit determines which controller is associated with the value of -1 in the **controller** parameter of an NI-VXI function. This new option indicates whether the -1 value is referring to the PCI-MXI-2 (the local controller) or the first remote VXI/VME-MXI-2 controller. For more details, refer to the *Default Controller (LA -1)* section in Chapter 6, *NI-VXI Configuration Utility*, of your getting started manual.

Remote controllers, when configured to detect asynchronous events such as a VXI interrupt or VXI trigger, need to inform the local controller that such an asynchronous event has occurred. The remote controllers report these events back to the local controller via a VXI IRQ line. This IRQ line is called the *system IRQ line*. You can use the NI-VXI utility VXIedit to select which VXI interrupt line the remote controller uses to report remote events to the local controller. You need to map the system IRQ line back to the local controller to receive remote controller interrupts. This mapping is performed automatically by the Resource Manager in the parent-side VXI-MXI-2 controllers, but not in other mainframe extenders. You can map interrupts by using the **Interrupt Configuration Editor** in VXIedit, or by using the `MapVXIInt()` function, which is described in Chapter 13, *VXIbus Extender Functions*, in the *NI-VXI Software Reference Manual for C*.

Notice the following differences in how the system IRQ line is treated differently than other IRQ lines used by NI-VXI.

- The system IRQ line is always acknowledged by the Resource Manager (Logical Address 0).
- The system IRQ line cannot be disabled on the Resource Manager. Calling `DisableVXIInt()` on the system IRQ line does *not* disable it.
- Devices other than remote controllers can also interrupt on the system IRQ line, provided that the device at Logical Address 0 is the handler for the interrupt.
- Routing the system IRQ to the signal queue is not recommended. Because the system IRQ cannot be disabled, it is possible that this routing could cause interrupts to be lost.

Enhancements to the NI-VXI Software

The NI-VXI software for the PCI-MXI-2 makes use of many new features available in MXI-2 and on the MITE ASIC. You have two options for controlling these features—through new configuration options in VXIedit, or by extensions to the NI-VXI Application Programming Interface (API). Refer to Chapter 6, *NI-VXI Configuration Utility*, and Chapter 7, *Using the NI-VXI Software*, in your getting started manual for more information on the new configuration options.

NI-VXI API Extensions

The following paragraphs discuss the additional options beyond what was documented in the *NI-VXI Software Reference Manual for C*.

Compatibility

NI-VXI applications that follow the guidelines in the *NI-VXI Software Reference Manual for C* will work with NI-VXI for the PCI-MXI-2 with a few exceptions. These exceptions deal primarily with low-level VXIbus access functions.

- ◆ **DOS/Windows 3.1 users**—Applications that use `VXIpeek()` and `VXIpoke()` may need to be recompiled. The `VXIpeek()` and `VXIpoke()` macros in other hardware platforms accessed locations in I/O space to check for retries on the MXIbus. PCI-MXI-2 registers are no longer located in I/O space, and the PCI-MXI-2 automatically handles these retry situations, so the macros have changed. Previous releases of NI-VXI provided a flag, `MXI1AND2_COMPATIBLE`. Applications compiled with this flag use PCI-MXI-2 binary-compatible `VXIpeek()` and `VXIpoke()` macros and will not require recompilation.

Applications that set the context bits directly for use in `SetContext()` may not be compatible with the new format for context. Because the PCI-MXI-2 allows more flexible window mapping, extra bits have been added to this field to reflect these new features. In general, we recommend that you do not manipulate the context bits.

System Configuration Functions

The **mainframe extender/controller information** field of the `DevInfo` structure (field 17) defines a new bit. Bit 13 reports if the extender is a remote controller. This bit is now significant in any function that uses the **mainframe extender/controller information** field (`GetDevInfo()`, `SetDevInfo()`, `GetDevInfoShort()`, `SetDevInfoShort()`, and `FindDevLA()`). NI-VXI automatically sets bit 13. Changing the value of this bit in a program, even via the documented API calls, is not recommended.

If a remote controller is sharing its onboard RAM, you can find the size, space, and offset of that RAM in the VXI system by calling `GetDevInfo()` for the remote controller's logical address.

Low-Level VXIbus Access Functions

Do not make any assumptions about the size and features of a window returned from `MapVXIAddress()`. You should use `GetWindowRange()` to determine the size of a window.

The 32-bit value returned from `GetContext()` and passed to `SetContext()` has a new format.

A new function, `MapVXIAddressSize()`, has been added to the low-level VXIbus access functions. Its description follows.

MapVXIAddressSize

`MapVXIAddressSize()` is the interface for requesting a particular window size for all successive calls to `MapVXIAddress()`.

Syntax: `ret = MapVXIAddressSize(size)`

Action: Sets the default size for the window returned from `MapVXIAddress()`.

Remarks: Input Parameter:

size	uint32	Size in bytes of window to be mapped
-------------	--------	--------------------------------------

Output Parameters:

none

Return Value:

ret	int16	Return Status 0 = Successful
------------	-------	---------------------------------

Example:

```
/* Set the size of future mappings to 1 MB */
```

```
ret = MapVXIAddressSize(0x100000);
```

High-Level VXIbus Access Functions

As previously mentioned, `VXImove()` has new bits defined in the **srcParm** and **destParm** fields for using features of the MITE DMA channel.

- Bit 11 — Use programmed I/O (PIO) instead of DMA. DMA is used by default and setting this bit forces NI-VXI to use PIO instead of DMA. In general, DMA is more efficient than PIO. However, when virtual memory is involved, a DMA channel needs to lock down memory before it can be used (locking is done by the NI-VXI DMA manager automatically). Notice that virtual memory is enabled under Windows 95/NT/3.1, but it is not available under DOS. Because of the amount of time required for buffer locking, it may be more efficient to transfer buffers by PIO.

- Bit 12 — No increment. Do not increment the address. This is used when the source or destination is a FIFO buffer and is valid for any address space, whether local, A16, A24, or A32.
- Bit 13 — Disable MXI block mode. MXI block mode is enabled by default in `VXImove()` when `VXImove()` is called from a CPU-MXI because it makes MXI transfers more efficient, and MXI block mode is supported by the VXI/VME-MXI-2. However, if you are communicating with a MXI device that does not support MXI block mode, you need to disable this mode by setting this bit.
- Bit 14 — Disable MXI synchronous mode. MXI synchronous mode is enabled by default in `VXImove()` when `VXImove()` is called from a CPU-MXI because it makes MXI transfers more efficient, and MXI synchronous mode is supported by the VXI/VME-MXI-2. However, if you are communicating with a MXI device that does not support MXI synchronous mode, you need to disable this mode by setting this bit.

In addition to the above bits, two new access privileges have been added to the **srcParm** and **destParm** fields to support the VME64 protocol.

- Access Privilege 6 — VME64 nonprivileged block access.
- Access Privilege 7 — VME64 privileged block access.

Setting these access privileges will cause `VXImove()` to perform the VME64 protocol on the VXI bus. Ensure that the device you are communicating with is capable of responding to these special cycles before using these privileges.

For best performance, keep the following in mind when using `VXImove()`.

- Make sure your buffers are 32-bit aligned.
- Transfer 32-bit data whenever possible.
- `VXImove()` must lock the user buffer in memory on virtual memory systems, so locking the buffer yourself will optimize `VXImove()`.
- Because `VXImove()` must build a scatter-gather list for the user buffer on paged memory systems, using a contiguous buffer will optimize `VXImove()`.
- Using VXI block access privileges will significantly improve performance to devices that are capable of accepting block transfers.
- `VXImemAlloc()` will return 32-bit aligned, page-locked, contiguous buffers, which work efficiently with `VXImove()`.

Local Resource Access Functions

`VXIMemAlloc()` does not allocate onboard RAM on the PCI-MXI-2; it only allocates system RAM on the motherboard. If you want to access onboard RAM on the PCI-MXI-2, access it as if it were VXI memory—that is, by using high-level or low-level VXIbus access functions. You can use `GetDevInfo()` on the PCI-MXI-2 device to determine the VXI address space and VXI address of this onboard RAM.

VXI Interrupt Functions

Every function that takes a controller's logical address as an argument can now accept the local controller or a remote controller's logical address as well.

VXI Trigger Functions

Every function that takes a controller's logical address as an argument can now accept the local controller or a remote controller's logical address as well.

TIC functionality is now available on the local controller and on every remote controller.

System Interrupt Handler Functions

Every function that takes a controller's logical address as an argument can now accept the local controller or a remote controller's logical address as well.