

## CONVERTING A MESSAGE SIGNED INTERRUPTION INTO AN I/O ADAPTER EVENT NOTIFICATION

### BACKGROUND

5 This invention relates, in general, to interruption processing within a computing environment, and in particular, to handling interruptions generated by adapters of the computing environment.

10 Message signaled interruption (MSI) is a way for an adapter function, such as a Peripheral Component Interconnect (PCI) function, to generate a central processing unit (CPU) interruption to notify the operating system of the occurrence of an event or the presence of some status. MSI is an alternative to having a dedicated interruption pin on each device. When an adapter function is configured to use MSI, the function requests an interruption by  
15 performing an MSI write operation of a specified number of bytes of data to a special address. The combination of this special address and a unique data value is termed an MSI vector.

20 Some adapter functions support only one MSI vector; other adapter functions support multiple MSI vectors. For functions that support multiple MSI vectors, the same special address is used with different data values.

25 In many computing platforms, a device driver configures itself as the interrupt handler associated with an MSI vector. This effectively associates an MSI vector with an entry in a CPU interruption vector. Therefore, when an adapter function supports multiple MSI vectors and is configured to use multiple MSI vectors, it consumes a corresponding number of entries in the CPU interruption vector.

30 US Publication No. 2007/0271559 A1, published November 22, 2007, Easton et al., "Virtualization of Infiniband Host Channel Adapter Interruptions," describes a method, system, program product and computer data structure for providing for two levels of server virtualization. A first hypervisor enables multiple logical partitions to share a set of

resources and provides a first level of virtualization. A second hypervisor enables multiple, independent virtual machines to share the resources that are assigned to a single logical partition and provides a second level of virtualization. All events for all of the virtual machines within said single logical partition are grouped into a single partition-owned event queue for receiving event notifications from the shared resources for that single logical partition. A request for an interrupt is signaled for the grouped events from the partition-owned event queue for the demultiplexing of grouped events, by the machine, from the partition-owned event queue into individual, virtualized event queues that are allocated on a per virtual machine basis.

US Publication No. 2005/0289271 A1, published December 29, 2005, Martinez, et al., "Circuitry to Selectively Produce MSI Signals," describes, in some embodiments, the inventions include a chip having a status register circuit coupled to conductors to receive interrupt event signals to provide source signals corresponding to the interrupt event signals. The chip also includes a control register circuit to provide source enable signals for selective ones of the interrupt sources, and a re-arming logic circuit coupled to the conductors to receive the interrupt event signals and provide a re-arming signal. The chip further includes first logic circuit to receive the source signals, the source enable signals, and the re-arming signal to provide an initial interrupt signal, and message signaled interrupt (MSI) signal pulse generation logic to receive the initial interrupt signal and provide an MSI signal in response thereto. Other embodiments are described and claimed.

US Patent No. 7,562,366, issued July 14, 2009, Pope et al., "Transmit Completion Event Batching," describes a method for managing a data transmit queue, for use with a host and a network interface device. Roughly described, the host writes data buffer descriptors into a transmit descriptor queue, and the network interface device writes events to notify the host when it has completed processing of a transmit data buffer. Each of the transmit completion event descriptors notify the host of completion of a plurality of the transmit data buffers.

**BRIEF SUMMARY**

In accordance with an aspect of the present invention, a capability is provided to facilitate managing of interruption requests from adapters.

5 The shortcomings of the prior art are overcome and advantages are provided through the provision of a method according to claim 1, and corresponding system and computer program product for managing interruption requests in a computing environment.

**BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS**

10 One or more aspects of the present invention are particularly pointed out and distinctly claimed as examples in the claims at the conclusion of the specification. The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

15 FIG. 1 depicts one embodiment of a computing environment to incorporate and use one or more aspects of the present invention;

FIG. 2 depicts one embodiment of further details of system memory and the I/O hub  
20 of FIG. 1, in accordance with an aspect of the present invention;

FIGs. 3A-3B depict examples of allocations of adapter interruption bit vectors, in accordance with an aspect of the present invention;

25 FIGs. 3C-3D depict examples of allocations of adapter interruption summary bits, in accordance with an aspect of the present invention;

FIG. 4 depicts one embodiment of an overview of the logic to be performed during initialization to configure an adapter function for I/O adapter event notification, in  
30 accordance with an aspect of the present invention;

FIG. 5 depicts one embodiment of the logic to perform registration to enable conversion of a message signaled interruption (MSI) into an I/O adapter event notification, in accordance with an aspect of the present invention;

5           FIG. 6A depicts one embodiment of the logic to convert an MSI request to an I/O adapter event notification, in accordance with an aspect of the present invention;

10           FIG. 6B depicts one embodiment of the logic to present the I/O adapter event notification to an operating system, in accordance with an aspect of the present invention;

FIG. 7A depicts one embodiment of a Modify PCI Function Controls instruction used in accordance with an aspect of the present invention;

15           FIG. 7B depicts one embodiment of a field used by the Modify PCI Function Controls instruction of FIG. 7A, in accordance with an aspect of the present invention;

FIG. 7C depicts one embodiment of another field used by the Modify PCI Function Controls instruction of FIG. 7A, in accordance with an aspect of the present invention;

20           FIG. 7D depicts one embodiment of the contents of a function information block (FIB) used in accordance with an aspect of the present invention;

25           FIG. 8 depicts one embodiment of an overview of the logic of the Modify PCI Function Controls instruction, in accordance with an aspect of the present invention;

FIG. 9 depicts one embodiment of the logic associated with a register adapter interruptions operation that may be specified by the Modify PCI Function Controls instruction, in accordance with an aspect of the present invention;

30           FIG. 10 depicts one embodiment of the logic associated with an unregister adapter interruptions operation that may be specified by the Modify PCI Function Controls instruction, in accordance with an aspect of the present invention;

FIG. 11A depicts one embodiment of a Call Logical Processor instruction used in accordance with an aspect of the present invention;

FIG. 11B depicts one embodiment of a request block used by the Call Logical Processor instruction of FIG. 11A for a list operation, in accordance with an aspect of the present invention;

FIG. 11C depicts one embodiment of a response block for the list operation of FIG. 11B, in accordance with an aspect of the present invention;

FIG. 11D depicts one embodiment of a function list entry used in accordance with an aspect of the present invention;

FIG. 12A depicts one embodiment of a request block used by the Call Logical Processor instruction of FIG. 11A for a query function operation, in accordance with an aspect of the present invention;

FIG. 12B depicts one embodiment of a response block for the query function operation of FIG. 12A, in accordance with an aspect of the present invention;

FIG. 13A depicts one embodiment of a request block used by the Call Logical Processor instruction of FIG. 11A for a query group operation, in accordance with an aspect of the present invention;

FIG. 13B depicts one embodiment of a response block for the query group operation of FIG. 13A, in accordance with an aspect of the present invention;

FIG. 14 depicts one embodiment of a computer program product incorporating one or more aspects of the present invention;

FIG. 15 depicts one embodiment of a host computer system to incorporate and use one or more aspects of the present invention;

FIG. 16 depicts a further example of a computer system to incorporate and use one or more aspects of the present invention;

FIG. 17 depicts another example of a computer system comprising a computer network to incorporate and use one or more aspects of the present invention;

FIG. 18 depicts one embodiment of various elements of a computer system to incorporate and use one or more aspects of the present invention;

FIG. 19A depicts one embodiment of the execution unit of the computer system of FIG. 18 to incorporate and use one or more aspects of the present invention;

FIG. 19B depicts one embodiment of the branch unit of the computer system of FIG. 18 to incorporate and use one or more aspects of the present invention;

FIG. 19C depicts one embodiment of the load/store unit of the computer system of FIG. 18 to incorporate and use one or more aspects of the present invention; and

FIG. 20 depicts one embodiment of an emulated host computer system to incorporate and use one or more aspects of the present invention.

## DETAILED DESCRIPTION

In accordance with an aspect of the present invention, a capability is provided for converting a message signaled interruption (MSI) request into an input/output (I/O) adapter event notification. The MSI is requested by an adapter and converted to an adapter event notification, in which one or more specific indicators are set and a request is made that an interruption be presented to an operating system (or other software, such as other programs, etc. As used herein, the term operating system includes operating system device drivers). In

one particular example, each MSI request does not result in a request for interruption to the operating system, but instead, one interruption request encompasses a plurality of MSI requests.

5 As used herein, the term "adapter" includes any type of adapter (e.g., storage adapter, network adapter, processing adapter, cryptographic adapter, PCI adapter, other type of input/output adapter, etc.). In one embodiment, an adapter includes one adapter function. However, in other embodiments, an adapter may include a plurality of adapter functions. One or more aspects of the present invention are applicable whether an adapter includes one  
10 adapter function or a plurality of adapter functions. Further, in the examples presented herein, adapter is used interchangeably with adapter function (e.g., PCI function) unless otherwise noted.

One embodiment of a computing environment to incorporate and use one or more aspects of  
15 the present invention is described with reference to FIG. 1. In one example, a computing environment 100 is a System z<sup>®</sup> server offered by International Business Machines Corporation. System z<sup>®</sup> is based on the z/Architecture<sup>®</sup> offered by International Business Machines Corporation. Details regarding the z/Architecture<sup>®</sup> are described in an IBM<sup>®</sup> publication entitled, "z/Architecture Principles of Operation," IBM Publication No. SA22-  
20 7832-07, February 2009. IBM<sup>®</sup>, System z<sup>®</sup> and z/Architecture<sup>®</sup> are registered trademarks of International Business Machines Corporation, Armonk, New York. Other names used herein may be registered trademarks, trademarks or product names of International Business Machines Corporation or other companies.

25 In one example, computing environment 100 includes one or more central processing units (CPUs) 102 coupled to a system memory 104 (a.k.a., main memory) via a memory controller 106. To access system memory 104, a central processing unit 102 issues a read or write request that includes an address used to access system memory. The address included in the request is typically not directly usable to access system memory, and therefore, it is  
30 translated to an address that is directly usable in accessing system memory. The address is translated via a translation mechanism (XLATE) 108. For example, the address is translated

from a virtual address to a real or absolute address using, for instance, dynamic address translation (DAT).

5 The request, including the address (translated, if necessary), is received by memory controller 106. In one example, memory controller 106 is comprised of hardware and is used to arbitrate for access to the system memory and to maintain the memory's consistency. This arbitration is performed for requests received from CPUs 102, as well as for requests received from one or more adapters 110. Like the central processing units, the adapters issue requests to system memory 104 to gain access to the system memory.

10 In one example, adapter 110 is a Peripheral Component Interconnect (PCI) or PCI Express (PCIe) adapter that includes one or more PCI functions. A PCI function issues a request that is routed to an input/output hub 112 (e.g., a PCI hub) via one or more switches (e.g., PCIe switches) 114. In one example, the input/output hub is comprised of hardware, including one or more state machines.

15 The input/output hub includes, for instance, a root complex 116 that receives the request from a switch. The request includes an input/output address that is used to perform a direct memory access (DMA) or to request a message signaled interruption (MSI), as examples.

20 This address is provided to an address translation and protection unit 118 which accesses information used for either the DMA or the MSI request.

25 For a DMA operation, address translation and protection unit 118 may translate the address to an address usable to access system memory. Then, the request initiated from the adapter, including the translated address, is provided to memory controller 106 via, for instance, an I/O-to-memory bus 120. The memory controller performs its arbitration and forwards the request with the translated address to the system memory at the appropriate time.

30 For an MSI request, information in address translation and protection unit 118 is obtained to facilitate conversion of the MSI request to an I/O adapter event notification. Since the embodiments described herein relate to interrupt processing, further details regarding the I/O hub and system memory as they relate to interrupt processing are described with reference to

FIG. 2. In FIG. 2, the memory controller is not shown, but may be used. The I/O hub may be coupled to system memory 104 and/or processor 254 directly or via a memory controller.

Referring to FIG. 2, in one example, system memory 104 includes one or more data structures usable in facilitating interruption processing. In this example, system memory 104 includes an adapter interruption bit vector (AIBV) 200 and an optional adapter interruption summary bit (AISB) 202 associated with a particular adapter. There may be an AIBV and a corresponding AISB for each adapter.

In one example, adapter interruption bit vector 200 is a single dimension array of one or more bits in main storage that are associated with an adapter (e.g., a PCI function). The bits in the adapter interruption bit vector represent MSI vector numbers. A bit that is set to one in an AIBV indicates a condition or type of event for the associated adapter. In the example of a PCI function, each bit in the associated AIBV corresponds to an MSI vector. Therefore, if a PCI function supports only one MSI vector, its AIBV includes a single bit; if a PCI function supports multiple MSI vectors, its AIBV includes one bit per MSI vector. In the example depicted in FIG. 2, the PCI function supports multiple MSI vectors (e.g., 3), and therefore, there are multiple bits (e.g., 3) in AIBV 200. Each bit corresponds to a particular event, e.g., bit 0 of the AIBV, when set to one, indicates a completed operation; bit 1 of the AIBV, when set to one, corresponds to an error event; etc. As shown, bit 1 is set in this example.

In one particular example, a command (e.g., a Modify PCI Function Controls command) is used to designate an AIBV for a PCI function. Specifically, the command is issued by the operating system and specifies the identity of the PCI function, the main storage location of the area that includes the AIBV, the offset from that location to the first bit of the AIBV, and the number of bits that comprise the AIBV. In particular, using this command, adapter interruption parameters are copied from a function information block that stores such information (e.g., obtained from initialization and/or configuration) into the adapter's device table entry (described below) and/or function table entry (described below).

The identity of the PCI function, in one example, is a function handle. A function handle includes, for instance, an enable indicator that indicates whether the PCI function handle is enabled; a PCI function number that identifies the function (this is a static identifier); and an instance number which indicates the particular instance of this function handle. For instance, each time the function handle is enabled, the instance number is incremented to provide a new instance number. The function handle is used to locate a function table entry in a function table that includes one or more entries. For instance, one or more bits of the function handle are used as an index into the function table to locate a particular function table entry. The function table entry includes information regarding its associated PCI function. For example, it may include various indicators regarding the status of its associated adapter function, and it may include one or more device table entry indices used to locate device table entries for this adapter function. (To the operating system, in one embodiment, the handle is simply an opaque identifier of the adapter.)

An AIBV may be allocated on any byte boundary and any bit boundary. This allows the operating system the flexibility to pack the AIBVs of multiple adapters into a contiguous range of bits and bytes. For instance, as shown in FIG. 3A, in one example, the operating system has designated a common storage area at location X to include five contiguous AIBVs. The adapter associated with each AIBV is identified by the letters A-E. The event that each AIBV bit represents for an adapter is further identified by the numbers 0-n. Unassigned bits are identified by the lowercase letter "u".

A further example is depicted in FIG. 3B. In this example, the operating system has designated three unique storage areas, at locations X, Y and Z to include the AIBVs for five I/O adapters. The storage at location X includes the AIBVs for adapters A and B, the storage at location Y includes the AIBV for only adapter C, and the storage at location Z includes the AIBVs for adapters D and E. The event that each AIBV bit represents for an I/O adapter is further identified by the numbers 0-n. Unassigned bits are identified by the letter "u".

Returning to FIG. 2, in addition to the AIBV, in this example, there is an AISB 202 for the adapter, which includes a single bit associated with the adapter. An AISB that is one

indicates that one or more bits have been set to one in an AIBV associated with the AISB. The AISB is optional, and there may be one for each adapter, one for each selected adapter or one for a group of adapters.

5 In one particular implementation for PCI functions, a command (e.g., a Modify PCI Function Controls command) is used to designate an AISB for a PCI function. Specifically, the command is issued by the operating system and specifies the identity of the PCI function (e.g., the handle), the main storage location of the area that includes the AISB, the offset from that location to the AISB, and an adapter interruption summary notification enablement  
10 control indicating there is a summary bit.

An AISB may be allocated on any byte boundary and any bit boundary. This allows the operating system the flexibility to pack the AISBs of multiple adapters into a contiguous range of bits and bytes. In one example, as depicted in FIG. 3C, the operating system has  
15 designated a common storage area, at location X, to include nine contiguous AISBs. The adapter associated with each AISB is identified by the letters A-I. Unassigned bits are identified by the lowercase letter "u".

A further allocation example is depicted in FIG. 3D, where the operating system has  
20 designated three unique AISB storage locations, at locations X, Y and Z to include the AISBs of each of the three adapters. The adapters associated with each AISB is identified with the letters A-C. Unassigned bits are identified by the lowercase letter "u".

Further, the program may also assign a single AISB to multiple PCI functions. This  
25 associates multiple AIBVs with a single summary bit. Therefore, such an AISB that is one indicates that the operating system should scan multiple AIBVs.

Returning to FIG. 2, in one example, the AIBV and the AISB are pointed to by addresses located in a device table entry 206 of a device table 208 located in I/O hub 112. In one  
30 example, device table 208 is located within the address translation protection unit of the I/O hub.

Device table 208 includes one or more entries 206, each of which is assigned to a particular adapter function 210. A device table entry 206 includes a number of fields, which may be populated using, for instance, the above-mentioned commands. The values of one or more of the fields are based on policy and/or configuration. Examples of the fields include:

5

Interruption Subclass (ISC) 214: Indicates an interruption subclass for the interruption. The ISC identifies a maskable class of adapter interruptions that may be associated with a priority with which the operating system will process the interruption;

10

AIBV Address (@) 216: Provides, e.g., an absolute address of the beginning of the storage location that includes the AIBV for the particular adapter function assigned to this device table entry;

15

AIBV Offset 218: An offset into the main storage location to the beginning of the AIBV;

AISB Address (@) 220: Provides, e.g., an absolute address of the beginning of the storage location that includes the AISB for this PCI function, if the operating system has designated an AISB;

20

AISB Offset 222: An offset into the main storage location to the AISB;

Adapter Interruption Summary Notification Enablement Control (Enable) 224: This control indicates whether there is an AISB;

25

Number of Interruptions (NOI) 226: Indicates the maximum number of MSI vectors allowed for this PCI function, with zero indicating none allowed.

In other embodiments, the DTE may include more, less or different information.

30

In one embodiment, the device table entry to be used for a particular interruption request by an adapter is located using, for instance, a requestor identifier (RID) (and/or a portion of the address) located in a request issued by the adapter (e.g., PCI function 210). The requestor

ID (e.g., a 16-bit value specifying, for instance, a bus number, device number and function number) is included in the request, as well as an address to be used for the interrupt. The request, including the RID and the address, are provided to, e.g., a contents addressable memory (CAM 230) via, e.g., a switch, and the contents addressable memory is used to provide an index value. For instance, the CAM includes multiple entries, with each entry corresponding to an index into the device table. Each CAM entry includes the value of a RID. If, for instance, the received RID matches the value contained in an entry in the CAM, the corresponding device table index is used to locate the device table entry. That is, the output of the CAM is used to index into device table 208. If there is no match, the received packet is discarded. (In other embodiments, a CAM or other lookup is not needed and the RID is used as the index.) The located DTE is used in processing an interrupt request, as described herein.

To request an interruption, adapter function 210 sends a packet to the I/O hub. This packet has an MSI address 232 and associated data 234. The I/O hub compares at least a part of the received address to a value in a MSI compare register 250. If there is a match, then an interruption (e.g., MSI) is being requested, as opposed to a DMA operation. The reason for the request (i.e., type of event that has occurred) is indicated in associated data 234. For example, one or more of the low order bits of the data are used to specify a particular interrupt vector (i.e., an MSI vector) that indicates the reason (event).

In accordance with an aspect of the present invention, the interruption request received from the adapter is converted into an I/O adapter event notification. That is, one or more indicators (e.g., one or more AIBVs and optionally an AISB) are set and an interruption to the operating system is requested, if one is not already pending. In one embodiment, multiple interruption requests (e.g., MSIs) from one or more adapters are coalesced into a single interruption to the operating system but with respective AIBV and AISB indications. For instance, if the I/O hub has already received an MSI request, has, in turn, provided an interruption request to a processor, and that interruption is still pending (e.g., for one reason or another, the interruption has not been presented to the operating system (e.g., interrupts are disabled)), then if the hub receives one or more other MSIs, it does not request additional

interrupts. The one interruption replaces and represents the plurality of MSI requests. However, one or more AIBVs and optionally one or more AISBs are set.

Further details regarding converting an MSI (or other adapter interruption request) to an I/O adapter event notification are described below with reference to FIGs. 4-6B. Particularly, FIG. 4 describes various initialization to be performed; FIG. 5 describes a registration process; FIG. 6A describes logic to convert an MSI to an adapter event notification; and FIG. 6B describes logic to present the I/O adapter event notification to the operating system.

Referring to FIG. 4, in one example, to convert an MSI request to an I/O adapter event notification certain initialization is performed. During initialization, the operating system performs a number of steps to configure an adapter for adapter event notification via an MSI request. In this example, it is a PCI function being configured; although, in other embodiments, it can be other adapters, including other types of adapter functions.

Initially, in one embodiment, a determination is made as to the PCI functions in the configuration, STEP 400. In one example, a command (e.g., a Query List command) issued by the operating system is used to obtain a list of the PCI functions assigned to the requesting configuration (e.g., assigned to a particular operating system). This information is obtained from a configuration data structure that maintains this information.

Next, one of the PCI functions in the list is selected, STEP 402, and a determination is made as to the MSI address to be used for the PCI function and the number of MSI vectors supported by the PCI function. The MSI address is determined based on the characteristics of the I/O hub and the system in which it is installed. The number of MSI vectors supported is based on policy and is configurable.

Additionally, the AIBV is allocated, as well as the AISB, if any, STEP 410. In one example, the operating system determines the location of the AIBV to allow for efficient processing of one or more adapters, typically based on the class of adapter. For example, the AIBVs for storage adapters may be located adjacent to each other. The AIBV and the AISB are allocated and cleared to zeros, and a register adapter interruption operation is specified (e.g.,

using a Modify PCI Function Controls instruction). This operation registers the AIBV, the AISB, the ISC, the number of interruptions (MSI vectors), and the adapter interruption summary notification enablement control, as described in further detail below, STEP 412. Thereafter, the PCI function's configuration space is read/written, STEP 414. Specifically,  
5 the MSI address and MSI vector count are written consistent with the previous registration.

Thereafter, a determination is made as to whether there are additional functions in the list, INQUIRY 416. If so, processing continues with STEP 402. Otherwise, initialization processing is complete.

10 Further details regarding the registration of various parameters are described with reference to FIG. 5. Initially, the device table entry (DTE) to correspond to the PCI function for which initialization is being performed is selected. This selection is performed by, for instance, the managing firmware that selects an available DTE from the device table. Thereafter, the  
15 various parameters are stored in the device table entry, STEP 502. For instance, the ISC, the AIBV address, the AIBV offset, the AISB address, the AISB offset, the enablement control, and the number of interruptions (NOI) are set to values obtained from configuring the function. This completes the registration process.

20 As used herein, firmware includes, e.g., the microcode, millicode and/or macrocode of the processor. It includes, for instance, the hardware-level instructions and/or data structures used in implementation of higher-level machine code. In one embodiment, it includes, for instance, proprietary code that is typically delivered as microcode that includes trusted software or microcode specific to the underlying hardware and controls operating system  
25 access to the system hardware.

During operation, when a PCI function wishes to generate an MSI, it typically makes some information available to the operating system that describes the condition. This causes one or more steps to occur in order to convert the PCI function's MSI request to an I/O adapter event notification to the operating system. This is described with reference to FIG. 6A.  
30

Referring to FIG. 6A, initially, a description of the event for which the interruption is requested is recorded, STEP 600. For instance, the PCI function records a description of the event in one or more adapter-specific event-description-recording structures stored, for instance, in system memory. This may include recording the type of event, as well as recording additional information. Additionally, a request is initiated by the PCI function specifying the MSI address and the MSI vector number, as well as a requestor ID, STEP 601. This request is received by the I/O hub, and responsive to receiving the request, the requestor ID in the request is used to locate the device table entry for the PCI function, STEP 602. The I/O hub compares at least a portion of the address in the request with the value in the MSI compare register, INQUIRY 603. If they are unequal, an MSI is not being requested. However, if they are equal, then an MSI address has been specified, and thus, an MSI has been requested, instead of a direct memory access operation.

Thereafter, a determination is made as to whether the MSI vector number specified in the request is less than or equal to the number of interruptions (NOI) allowed for this function, INQUIRY 604. If the MSI vector number is greater than NOI, an error is indicated. Otherwise, the I/O hub issues a set bit function to set the appropriate AIBV bit in storage. The appropriate bit is determined by adding the MSI vector number to the AIBV offset specified in the device table entry and displacing this number of bits from the AIBV address specified in the device table entry, STEP 605. Moreover, if an AISB has been designated, the I/O hub uses a set bit function to set the AISB, using the AISB address and the AISB offset in the device table entry, STEP 606.

Next, in one embodiment, a determination is made (e.g., by the CPU or the I/O hub) as to whether an interruption request is already pending. To make this determination, a pending indicator is used. For instance, a pending indicator 252 (FIG. 2) stored in memory of a processor 254, which is accessible to processors of the computing environment that may process the interrupt (e.g., CPUs 102 of FIG. 1), is checked, INQUIRY 608. If it is not set, then it is set (e.g., to 1), STEP 610. If it is already set, processing is complete and another interruption request is not requested. Therefore, subsequent interruption requests are encompassed by the one request already pending.

In one particular example, there may be one pending indicator per interruption subclass, and therefore, the pending indicator of the interruption subclass assigned to the requesting function is the indicator that is checked.

5 Asynchronously, as depicted in FIG. 6B, one or more processors check the pending indicator, INQUIRY 640. In particular, each processor enabled for the ISC (and zone in another embodiment) polls on the indicator for the ISC when, for instance, interrupts are enabled for that processor (i.e., for its operating system). If one of the processors determines that the indicator is set, it arbitrates with the other processors enabled for the same ISC (and  
10 zone in another embodiment) to present the interruption, STEP 642. Returning to INQUIRY 640, if the pending indicator is not set, the processors enabled for the ISC continue to poll for a set indicator.

Responsive to the operating system being presented with the interruption, STEP 642, the  
15 operating system determines whether any AISBs are registered, INQUIRY 643. If not, the operating system processes the set AIBVs, as described below, STEP 645. Otherwise, the operating system processes any set AISBs and AIBVs, STEPs 644, 645. For example, it checks whether any AISBs are set. If so, it uses the AISB to determine the location of one or more AIBVs. For example, the operating system remembers the locations of the AISBs and  
20 AIBVs. Furthermore, it remembers for which adapter each AISB and AIBV represents. Therefore, it may maintain a form of a control block or other data structure that includes the locations of AISBs and AIBVs and the association between AISBs, AIBVs and adapter ID. It uses this control block to facilitate the location of an AIBV based on its associated AISB. In a further embodiment, an AISB is not used. In that situation, the control block is used to  
25 locate the particular AIBV.

Responsive to locating the one or more AIBVs, the operating system scans the AIBVs and processes any set AIBVs. It processes the interruption in a manner consistent with the presented event (e.g., provides status). For example, with a storage adapter, an event may  
30 indicate that an operation has completed. This results in the operating system checking status stored by the adapter to see if the operation completed successfully and also details of

the operation. In the case of a storage read, this is an indication that the data read from the adapter is now available in system memory and can be processed.

In one embodiment, if during operation of the conversion, an error is detected, an attention is generated to the system firmware, instead of converting the MSI request to an adapter event notification.

Further details regarding the Modify PCI Function Controls instruction used to register adapter interruptions is described herein. Referring to FIG. 7A, a Modify PCI Function Controls instruction 700 includes, for instance, an op code 702 indicating the Modify PCI Function Controls instruction; a first field 704 specifying a location at which various information is included regarding the adapter function for which the operational parameters are being established; and a second field 706 specifying a location from which a PCI function information block (FIB) is fetched. The contents of the locations designated by Fields 1 and 2 are further described below.

In one embodiment, Field 1 designates a general register that includes various information. As shown in FIG. 7B, the contents of the register include, for instance, a function handle 710 that identifies the handle of the adapter function on behalf of which the modify instruction is being performed; an address space 712 designating an address space in system memory associated with the adapter function designated by the function handle; an operation control 714 which specifies the operation to be performed for the adapter function; and status 716 which provides status regarding the instruction when the instruction completes with a predefined code.

In one embodiment, the function handle includes, for instance, an enable indicator indicating whether the handle is enabled, a function number that identifies an adapter function (this is a static identifier and may be used to index into a function table); and an instance number specifying the particular instance of this function handle. There is one function handle for each adapter function, and it is used to locate a function table entry (FTE) within the function table. Each function table entry includes operational parameters and/or other

information associated with its adapter function. As one example, a function table entry includes:

Instance Number: This field indicates a particular instance of the adapter function handle associated with the function table entry;

Device Table Entry (DTE) Index 1...n: There may be one or more device table indices, and each index is an index into a device table to locate a device table entry (DTE). There are one or more device table entries per adapter function, and each entry includes information associated with its adapter function, including information used to process requests of the adapter function (e.g., DMA requests, MSI requests) and information relating to requests associated with the adapter function (e.g., PCI instructions). Each device table entry is associated with one address space within system memory assigned to the adapter function. An adapter function may have one or more address spaces within system memory assigned to the adapter function.

Busy Indicator: This field indicates whether the adapter function is busy;

Permanent Error State Indicator: This field indicates whether the adapter function is in a permanent error state;

Recovery Initiated Indicator: This field indicates whether recovery has been initiated for the adapter function;

Permission Indicator: This field indicates whether the operating system trying to control the adapter function has authority to do so;

Enable Indicator: This field indicates whether the adapter function is enabled (e.g., 1=enabled, 0=disabled);

Requestor Identifier (RID): This is an identifier of the adapter function, and includes, for instance, a bus number, a device number and a function number.

In one example, this field is used for accesses of a configuration space of the adapter function. (Memory of an adapter may be defined as address spaces, including, for instance, a configuration space, an I/O space, and/or one or more memory spaces.) In one example, the configuration space may be accessed by specifying the configuration space in an instruction issued by the operating system (or other configuration) to the adapter function. Specified in the instruction is an offset into the configuration space and a function handle used to locate the appropriate function table entry that includes the RID. The firmware receives the instruction and determines it is for a configuration space. Therefore, it uses the RID to generate a request to the I/O hub, and the I/O hub creates a request to access the adapter. The location of the adapter function is based on the RID, and the offset specifies an offset into the configuration space of the adapter function.

Base Address Register (BAR) (1 to n): This field includes a plurality of unsigned integers, designated as  $\text{BAR}_0 - \text{BAR}_n$ , which are associated with the originally specified adapter function, and whose values are also stored in the base address registers associated with the adapter function. Each BAR specifies the starting address of a memory space or I/O space within the adapter function, and also indicates the type of address space, that is whether it is a 64 or 32 bit memory space, or a 32 bit I/O space, as examples.

In one example, it is used for accesses to memory space and/or I/O space of the adapter function. For instance, an offset provided in an instruction to access the adapter function is added to the value in the base address register associated with the address space designated in the instruction to obtain the address to be used to access the adapter function. The address space identifier provided in the instruction identifies the address space within the adapter function to be accessed and the corresponding BAR to be used;

Size 1...n: This field includes a plurality of unsigned integers, designated as  $\text{SIZE}_0 - \text{SIZE}_n$ . The value of a Size field, when non-zero, represents the size of each address space with each entry corresponding to a previously described BAR.

Further details regarding BAR and Size are described below.

1. When a BAR is not implemented for an adapter function, the BAR field and its corresponding size field are both stored as zeros.

2. When a BAR field represents either an I/O address space or a 32-bit memory address space, the corresponding size field is non-zero and represents the size of the address space.

3. When a BAR field represents a 64-bit memory address space,

a. The  $\text{BAR}_n$  field represents the least significant address bits.

b. The next consecutive  $\text{BAR}_{n+1}$  field represents the most significant address bits.

c. The corresponding  $\text{SIZE}_n$  field is non-zero and represents the size of the address space.

d. The corresponding  $\text{SIZE}_{n+1}$  field is not meaningful and is stored as zero.

Internal Routing Information: This information is used to perform particular routing to the adapter. It includes, for instance, node, processor chip, and hub addressing information, as examples.

Status Indication: This provides an indication of, for instance, whether load/store operations are blocked or the adapter is in the error state, as well as other indications.

In one example, the busy indicator, permanent error state indicator, and recovery initiated indicator are set based on monitoring performed by the firmware. Further, the permission indicator is set, for instance, based on policy; and the BAR information is based on configuration information discovered during a bus walk by the processor (e.g., firmware of the processor). Other fields may be set based on configuration, initialization, and/or events. In other embodiments, the function table entry may include more, less or different information. The information included may depend on the operations supported by or enabled for the adapter function.

Referring to FIG. 7C, in one example, Field 2 designates a logical address 720 of a PCI function information block (FIB), which includes information regarding an associated adapter function. The function information block is used to update a device table entry and/or function table entry (or other location) associated with the adapter function. The information is stored in the FIB during initialization and/or configuration of the adapter, and/or responsive to particular events.

Further details regarding a function information block (FIB) are described with reference to FIG. 7D. In one embodiment, a function information block 750 includes the following fields:

Format 751: This field specifies the format of the FIB.

Interception Control 752: This field is used to indicate whether guest execution of specific instructions by a pageable mode guest results in instruction interception;

Error Indication 754: This field includes the error state indication for direct memory access and adapter interruptions. When the bit is set (e.g., 1), one or more errors have been detected while performing direct memory access or adapter interruption for the adapter function;

Load/Store Blocked 756: This field indicates whether load/store operations are blocked;

PCI Function Valid 758: This field includes an enablement control for the adapter function. When the bit is set (e.g., 1), the adapter function is considered to be enabled for I/O operations;

Address Space Registered 760: This field includes a direct memory access enablement control for an adapter function. When the field is set (e.g., 1) direct memory access is enabled;

Page Size 761: This field indicates the size of the page or other unit of memory to be accessed by a DMA memory access;

5 PCI Base Address (PBA) 762: This field is a base address for an address space in system memory assigned to the adapter function. It represents the lowest virtual address that an adapter function is allowed to use for direct memory access to the specified DMA address space;

10 PCI Address Limit (PAL) 764: This field represents the highest virtual address that an adapter function is allowed to access within the specified DMA address space;

15 Input/Output Address Translation Pointer (IOAT) 766: The input/output address translation pointer designates the first of any translation tables used by a PCI virtual address translation, or it may directly designate the absolute address of a frame of storage that is the result of translation;

Interruption Subclass (ISC) 768: This field includes the interruption subclass used to present adapter interruptions for the adapter function;

20 Number of Interruptions (NOI) 770: This field designates the number of distinct interruption codes accepted for an adapter function. This field also defines the size, in bits, of the adapter interruption bit vector designated by an adapter interruption bit vector address and adapter interruption bit vector offset fields;

25 Adapter Interruption Bit Vector Address (AIBV) 772: This field specifies an address of the adapter interruption bit vector for the adapter function. This vector is used in interrupt processing;

30 Adapter Interruption Bit Vector Offset 774: This field specifies the offset of the first adapter interruption bit vector bit for the adapter function;

Adapter Interruption Summary Bit Address (AISB) 776: This field provides an address designating the adapter interruption summary bit, which is optionally used in interrupt processing;

5       Adapter Interruption Summary Bit Offset 778: This field provides the offset into the adapter interruption summary bit vector;

10       Function Measurement Block (FMB) Address 780: This field provides an address of a function measurement block used to collect measurements regarding the adapter function;

Function Measurement Block Key 782: This field includes an access key to access the function measurement block;

15       Summary Bit Notification Control 784: This field indicates whether there is a summary bit vector being used;

20       Instruction Authorization Token 786: This field is used to determine whether a pageable storage mode guest is authorized to execute PCI instructions without host intervention.

25       In one example, in the z/Architecture<sup>®</sup>, a pageable guest is interpretively executed via the Start Interpretive Execution (SIE) instruction, at level 2 of interpretation. For instance, the logical partition (LPAR) hypervisor executes the SIE instruction to begin the logical partition in physical, fixed memory. If z/VM<sup>®</sup> is the operating system in that logical partition, it issues the SIE instruction to execute its guests (virtual) machines in its V=V (virtual) storage. Therefore, the LPAR hypervisor uses level-1 SIE, and the z/VM<sup>®</sup> hypervisor uses level-2 SIE; and

30       Address Translation Format 787: This field indicates a selected format for address translation of the highest level translation table to be used in translation (e.g., segment table, region 3rd, etc).

The function information block designated in the Modify PCI Function Controls instruction is used to modify a selected device table entry, a function table entry and/or other firmware controls associated with the adapter function designated in the instruction. By modifying the device table entry, function table entry and/or other firmware controls, certain services are provided for the adapter. These services include, for instance, adapter interruptions; address translations; reset error state; reset load/store blocked; set function measurement parameters; and set interception control.

One embodiment of the logic associated with the Modify PCI Function Controls instruction is described with reference to FIG. 8. In one example, the instruction is issued by an operating system (or other configuration) and executed by the processor (e.g., firmware) executing the operating system. In the examples herein, the instruction and adapter functions are PCI based. However, in other examples, a different adapter architecture and corresponding instructions may be used.

In one example, the operating system provides the following operands to the instruction (e.g., in one or more registers designated by the instruction): the PCI function handle; the DMA address space identifier; an operation control; and an address of the function information block.

Referring to FIG. 8, initially, a determination is made as to whether the facility allowing for a Modify PCI Function Controls instruction is installed, INQUIRY 800. This determination is made by, for instance, checking an indicator stored in, for instance, a control block. If the facility is not installed, an exception condition is provided, STEP 802. Otherwise, a determination is made as to whether the instruction was issued by a pageable storage mode guest (or other guest), INQUIRY 804. If yes, the host operating system will emulate the operation for that guest, STEP 806.

Otherwise, a determination is made as to whether one or more of the operands are aligned, INQUIRY 808. For instance, a determination is made as to whether the address of the function information block is on a double word boundary. In one example, this is optional. If the operands are not aligned, then an exception condition is provided, STEP 810.

Otherwise, a determination is made as to whether the function information block is accessible, INQUIRY 812. If not, then an exception condition is provided, STEP 814.

Otherwise, a determination is made as to whether the handle provided in the operands of the Modify PCI Function Controls instruction is enabled, INQUIRY 816. In one example, this  
5 determination is made by checking an enable indicator in the handle. If the handle is not enabled, then an exception condition is provided, STEP 818.

If the handle is enabled, then the handle is used to locate a function table entry, STEP 820.

That is, at least a portion of the handle is used as an index into the function table to locate the  
10 function table entry corresponding to the adapter function for which operational parameters are to be established.

A determination is made as to whether the function table entry was found, INQUIRY 822. If not, then an exception condition is provided, STEP 824. Otherwise, if the configuration  
15 issuing the instruction is a guest, INQUIRY 826, then an exception condition (e.g., interception to the host) is provided, STEP 828. This inquiry may be ignored if the configuration is not a guest or other authorizations may be checked, if designated.

A determination is then made as to whether the function is enabled, INQUIRY 830. In one  
20 example, this determination is made by checking an enable indicator in the function table entry. If it is not enabled, then an exception condition is provided, STEP 832.

If the function is enabled, then a determination is made as to whether recovery is active, INQUIRY 834. If recovery is active as determined by a recovery indicator in the function  
25 table entry, then an exception condition is provided, STEP 836. However, if recovery is not active, then a further determination is made as to whether the function is busy, INQUIRY 838. This determination is made by checking the busy indicator in the function table entry. If the function is busy, then a busy condition is provided, STEP 840. With the busy condition, the instruction can be retried, instead of dropped.

If the function is not busy, then a further determination is made as to whether the function  
30 information block format is valid, INQUIRY 842. For instance, the format field of the FIB

is checked to determine if this format is supported by the system. If it is invalid, then an exception condition is provided, STEP 844. If the function information block format is valid, then a further determination is made as to whether the operation control specified in the operands of the instruction is valid, INQUIRY 846. That is, is the operation control one of the specified operation controls for this instruction. If it is invalid, then an exception condition is provided, STEP 848. However, if the operation control is valid, then processing continues with the specific operation control being specified.

In one example, the operation control is a register adapter interruptions operation, which is used for controlling adapter interruptions. Responsive to this operation control, the adapter function parameters relevant to adapter interruptions are set in the device table entry based on the appropriate contents of the function information block.

One embodiment of the logic associated with this operation is described with reference to FIG. 9. As one example, the operands for this operation, which are obtained from the function information block, include for instance: an interruption subclass (ISC); number of interruptions allowed (NOI); an adapter interruption bit vector offset (AIBVO); a summary notification (S); an adapter interruption summary bit vector offset (ABVSO); an adapter interruption bit vector (AIBV) address; and an adapter interruption summary bit vector (AISB) address.

Referring to FIG. 9, initially, a determination is made as to whether the number of interruptions (NOIs) specified in the FIB is greater than a model-dependent maximum, INQUIRY 900. If so, then an exception condition is provided, STEP 902. However, if the number of interruptions is not greater than the model-dependent maximum, then a further determination is made as to whether the number of interruptions added to the adapter interruption bit vector offset ( $\text{NOI} + \text{AIBVO}$ ) is greater than a model-dependent maximum, INQUIRY 904. If so, then an exception condition is provided, STEP 906. If the NOI plus the AIBVO is not greater than a model-dependent maximum, then a further determination is made as to whether the AIBV address plus the NOI spans a 4k boundary, INQUIRY 908. If it does span the 4k boundary, then an exception condition is provided, STEP 910. Otherwise, a determination is made as to whether sufficient resources are available for any

resources needed, STEP 912. If there are not sufficient resources, then an exception condition is provided, STEP 914.

Otherwise, a determination is made as to whether adapter interruptions are already registered  
5 for this function, STEP 916. In one embodiment, this would be determined by checking one  
or more of the parameters (e.g., in the DTE/FTE). In particular, parameters associated with  
interruptions, such as NOI, are checked. If the fields are populated, then the adapter is  
registered for interrupts. If the adapter is already registered, then an exception condition is  
provided, STEP 918. Otherwise, the interruption parameters are obtained from the FIB and  
10 placed in the device table entry and optionally, in the corresponding function table entry  
(FTE) (or other specified location). Also, an MSI enablement indicator is set in the DTE,  
STEP 920. That is, the PCI function parameters relevant to adapter interruption are set in  
the DTE and optionally, in the FTE based on the information retrieved from the function  
information block. These parameters include, for instance, the ISC, NOI, AIBVO, S,  
15 AIBVSO, AIBV address and the AISB address.

In addition to the above, another operation control that can be specified is an unregister  
adapter interruptions operation, an example of which is described with reference to FIG. 10.  
With this operation, the adapter function parameters relevant to adapter interruption are  
20 reset.

Referring to FIG. 10, initially, a determination is made as to whether the adapter specified by  
the function handle is registered for interrupts, INQUIRY 1000. If not, then an exception  
condition is provided, STEP 1002. Otherwise, the interruption parameters in the function  
25 table entry (or other location) and corresponding device table entry are set to zeros,  
INQUIRY 1004. In one example, these parameters include the ISC, NOI, AIBVO, S,  
AIBSO, AIBV address and AISB address.

As described above, in one embodiment, to obtain the information regarding an adapter  
30 function, a Call Logical Processor instruction is used. One embodiment of this instruction is  
depicted in FIG. 11A. As shown, in one example, a Call Logical Processor (CLP)  
instruction 1100 includes an operation code 1102 indicating that it is the Call Logical

Processor instruction; and an indication for a command 1104. In one example, this indication is an address of a request block that describes the command to be performed, and the information in the request block is dependent on the command. Examples of requests blocks and corresponding response blocks for various commands are described with  
5 reference to FIGs. 11B-13B.

Referring initially to FIG. 11B, a request block for a list PCI functions command is provided. The list PCI functions command is used to obtain a list of PCI functions that are assigned to the requesting configuration (e.g., the requesting operating system). A request  
10 block 1120 includes a number of parameters, such as, for instance:

Length field 1122: This field indicates the length of the request block;

Command Code 1124: This field indicates the list PCI functions command; and  
15

Resume Token 1126: This field is an integer that is used to either start a new list PCI functions command or resume a previous list PCI functions command, as described in further detail below.

When the resume token field in the command request block includes, for instance, a value of zero, a new list of PCI functions is requested. When the resume token field includes, for instance, a non-zero value, which was returned from a previous list PCI functions command, a continuation of a previous list of PCI functions is requested.  
20

Responsive to issuing and processing the Call Logical Processor instruction for a list PCI functions command, a response block is returned. One embodiment of the response block is depicted in FIG. 11C. In one example, a response block 1150 for a list PCI functions command includes:  
25

Length field 1152: This field indicates the length of the response block;  
30

Response Code 1154: This field indicates a status of the command;

PCI Function List 1156: This field indicates a list of one or more PCI functions available to the requesting operating system;

Resume Token 1158: This field indicates whether a continuation of a previous list of PCI functions is requested. In one example, when the resume token in the request block and the resume token in the response block are zero, all PCI functions assigned to the requesting configuration are represented in the PCI function list; if the resume token in the request block is zero and the resume token in the response block is not zero, additional PCI functions assigned to the request configuration may exist that have not been represented in the list; if the resume token in the request block is not zero and the resume token in the response block is zero, from the resume point, remaining PCI functions assigned to the requesting configuration are represented in the list; when both the resume tokens in the request and response block are not zero from the resume point, additional PCI functions assigned to the requesting configuration may exist that have not been represented in any associated PCI function list. The resume token remains valid for an indefinite period of time after being returned, but it may be invalid due to a variety of model dependent reasons, including system load elapse time;

Model Dependent Data 1160: This field includes data that depends on the system;

Number of PCI Functions 1162: This field indicates the maximum number of PCI functions supported by the facility; and

Entry Size 1164: This field indicates the size of each entry in the PCI function list.

Further details regarding the PCI function list are described with reference to FIG. 11D. In one example, the PCI function list includes a plurality of entries and each entry 1156 includes the following information, as an example:

Device ID 1170: This field indicates the I/O adapter associated with the corresponding PCI function;

Vendor ID 1172: This field identifies the manufacturer of the I/O adapter associated with the corresponding PCI function;

Function Identifier 1174: This field includes a persistent identifier of the PCI function;

Function Handle 1176: This field identifies a PCI function. The PCI function handle stored is a general handle when a specified bit of the handle is zero, and it is an enabled handle when that bit is one. If the PCI function is disabled, a general PCI function handle is stored. If the PCI function is enabled, an enabled PCI function handle is stored. A PCI function handle is not, in one example, persistent beyond an IPL, which differs from the PCI function ID, which is persistent and is set for the life of the I/O configuration definition; and

Configuration State 1178: This field indicates the state of the PCI function. When this indicator is, for instance, zero, the state is standby, and when, for instance, one, the state is configured. When in standby, the PCI function handle is the general PCI function handle, and when configured, it is either the general or enabled PCI function handle depending on whether the PCI function is enabled.

Subsequent to obtaining the list of adapter functions, information may be obtained regarding the attributes of a selected function as designated by a specified PCI function handle. This information may be obtained by issuing a CLP instruction with a query function command.

One embodiment of the request block for a query PCI function command is described with reference to FIG. 12A. In one example, request block 1200 includes, for instance:

Length field 1202: This field indicates the length of the request block;

Command Code 1204: This field indicates the query PCI function command; and

Function Handle 1206: This field includes the PCI function handle (e.g., general or enabled) that designates the PCI function to be queried.

Responsive to issuing the Call Logical Processor instruction for the query PCI function command, a response block is returned. One embodiment of the response block is depicted in FIG. 12B. In one example, a response block 1250 includes the following:

Length 1252: This field indicates the length of the response block;

Response Code 1254: This field indicates a status of the command;

Function Group ID 1256: This field indicates the PCI function group identifier. A PCI function group identifier is used to associate a group of PCI functions with a set of attributes (also referred to herein as characteristics). Each PCI function with the same PCI function group identifier has the same set of attributes;

Function ID 1258: The PCI function id is a persistent identifier of the PCI function originally specified by the PCI function handle and is set for the life of the I/O configuration definition;

Physical Channel Adapter 1260: This value represents a model dependent identification of the location of the physical I/O adapter which corresponds to the PCI function;

Base Address Registers (BARs) 1...n 1262: This field includes a plurality of unsigned integers, designated as  $\text{BAR}_0 - \text{BAR}_n$ , which are associated with the originally specified PCI function, and whose values are also stored in the base address registers associated with the PCI function. Each BAR specifies the starting address of a memory space or I/O space within the adapter, and also indicates the type of address space, that is whether it is a 64 or 32 bit memory space, or a 32 bit I/O space, as examples;

Size 1...n 1264: This field includes a plurality of unsigned integers, designated as SIZE<sub>0</sub> – SIZE<sub>n</sub>. The value of a Size field, when non-zero, represents the size of each address space with each entry corresponding to a previously described BAR.

5           Start Available DMA 1266: This field includes an address which indicates the beginning of a range of PCI addresses that are available for DMA operations;

          End Available DMA 1268: This field includes a value which indicates the end of a range of PCI addresses that are available for DMA operations.

10

In addition to obtaining attributes regarding the specific adapter function, attributes may also be obtained regarding the group that includes this function. These common attributes may be obtained from issuing a CLP instruction with a query PCI function group command. This command is used to obtain a set of characteristics that are supported for a group of one or  
15       more PCI functions designated by the specified PCI function group identifier. A PCI function group identifier is used to associate a group of PCI functions with the same set of characteristics. One embodiment of a request block for the query PCI function group command is described with reference to FIG. 13A. In one example, request block 1300 includes the following:

20

          Length field 1302: This field indicates the length of the request block;

          Command Code 1304: This field indicates the query PCI function group command;  
          and

25

          Function Group ID 1306: This field specifies the PCI function group identifier for which attributes are to be obtained.

30

Responsive to issuing and processing the Call Logical Processor instruction with a query PCI function group command, a response block is returned. One embodiment of the response block is depicted in FIG. 13B. In one example, a response block 1350 includes:

Length Field 1352: This field indicates the length of the response block;

Response Code 1354: This field indicates a status of the command;

5           Number of Interruptions 1356: This field indicates the maximum number of consecutive MSI vector numbers (i.e., interruption event indicators) that are supported by the PCI facility for each PCI function in the specified PCI function group. The possible valid values of the number of interruptions are in the range of zero to 2,048, in one example;

10           Version 1358: This field indicates the version of the PCI specification that is supported by the PCI facility to which the group of PCI functions designated by the specified PCI group identifier are attached;

15           Frame 1362: This field indicates the frame (or page) sizes supported for I/O address translation;

20           Measurement Block Update Interval 1364: This is a value indicating the approximate time interval (e.g., in milliseconds) at which the PCI function measurement block is updating;

DMA Address Space Mask 1366: This is a value used to indicate which bits in a PCI address are used to identify a DMA address space; and

25           MSI Address 1368: This is a value that is to be used for message signal interruption requests.

30           The query list and function commands described above retrieve information from, for instance, the function table. At initialization time, or after a hot plug of an adapter, firmware performs a bus walk to determine the location of the adapter and determines its basic characteristics. This information is stored by the firmware into the function table entry (FTE) for each adapter. Accessibility to the adapter is determined based on policy set by a

system administrator and is also set by firmware into the FTE. The query list and function commands can then retrieve this information and store it in their respective response blocks accessible to the operating system.

5 Further, the group information is based on a given system I/O infrastructure and the capabilities of the firmware and the I/O hub. This may be stored in the FTE or any other convenient location for later retrieval during the query processing. In particular, the query group command retrieves the information and stores it in its response block accessible to the operating system.

10 Described in detail above is a capability for converting a PCI message signal interruption into an I/O adapter event notification to an operating system. This provides for a low latency interruption request; a delivery of MSIs from a relatively large number of PCI functions to the operating system; and the retaining of the flavor of the MSI vector designation that fits  
15 the MSI into the adapter event notification architecture. It adapts to allow the I/O hub to connect to a relatively large number of PCI functions and eliminates the problem of each write to an MSI vector generating a unique interruption.

In the embodiments described herein, the adapters are PCI adapters. PCI, as used herein,  
20 refers to any adapters implemented according to a PCI-based specification as defined by the Peripheral Component Interconnect Special Interest Group (PCI-SIG) ([www.pcisig.com/home](http://www.pcisig.com/home)), including but not limited to, PCI or PCIe. In one particular example, the Peripheral Component Interconnect Express (PCIe) is a component level interconnect standard that defines a bi-directional communication protocol for transactions  
25 between I/O adapters and host systems. PCIe communications are encapsulated in packets according to the PCIe standard for transmission on a PCIe bus. Transactions originating at I/O adapters and ending at host systems are referred to as upbound transactions. Transactions originating at host systems and terminating at I/O adapters are referred to as downbound transactions. The PCIe topology is based on point-to-point unidirectional links  
30 that are paired (e.g., one upbound link, one downbound link) to form the PCIe bus. The PCIe standard is maintained and published by the PCI-SIG.

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system". Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain or store a program for use by or in connection with an instruction execution system, apparatus, or device.

Referring now to FIG. 14, in one example, a computer program product 1400 includes, for instance, one or more computer readable storage media 1402 to store computer readable program code means or logic 1404 thereon to provide and facilitate one or more aspects of the present invention.

Program code embodied on a computer readable medium may be transmitted using an appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language, such as Java, Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language, assembler or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to

produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

5 The flowchart and block diagrams in the figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical  
10 function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and  
15 combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

In addition to the above, one or more aspects of the present invention may be provided,  
20 offered, deployed, managed, serviced, etc. by a service provider who offers management of customer environments. For instance, the service provider can create, maintain, support, etc. computer code and/or a computer infrastructure that performs one or more aspects of the present invention for one or more customers. In return, the service provider may receive payment from the customer under a subscription and/or fee agreement, as examples.  
25 Additionally or alternatively, the service provider may receive payment from the sale of advertising content to one or more third parties.

In one aspect of the present invention, an application may be deployed for performing one or more aspects of the present invention. As one example, the deploying of an application  
30 comprises providing computer infrastructure operable to perform one or more aspects of the present invention.

As a further aspect of the present invention, a computing infrastructure may be deployed comprising integrating computer readable code into a computing system, in which the code in combination with the computing system is capable of performing one or more aspects of the present invention.

5

As yet a further aspect of the present invention, a process for integrating computing infrastructure comprising integrating computer readable code into a computer system may be provided. The computer system comprises a computer readable medium, in which the computer medium comprises one or more aspects of the present invention. The code in combination with the computer system is capable of performing one or more aspects of the present invention.

10

Although various embodiments are described above, these are only examples. For example, computing environments of other architectures can incorporate and use one or more aspects of the present invention. As examples, servers other than System z<sup>®</sup> servers, such as Power Systems servers or other servers offered by International Business Machines Corporation, or servers of other companies can include, use and/or benefit from one or more aspects of the present invention. Further, although in the example herein, the adapters and PCI hub are considered a part of the server, in other embodiments, they do not have to necessarily be considered a part of the server, but can simply be considered as being coupled to system memory and/or other components of a computing environment. The computing environment need not be a server. Further, although the adapters are PCI based, one or more aspects of the present invention are usable with other adapters or other I/O components. Adapter and PCI adapter are just examples. Further, one or more aspects of the present invention are applicable to interruption schemes other than PCI MSI. Yet further, although examples are described in which bits are set, in other embodiments, bytes or other type of indicators may be set. Moreover, the DTE may include more, less or different information. Many other variations are possible.

15

20

25

30

Further, other types of computing environments can benefit from one or more aspects of the present invention. As an example, a data processing system suitable for storing and/or executing program code is usable that includes at least two processors coupled directly or

indirectly to memory elements through a system bus. The memory elements include, for instance, local memory employed during actual execution of the program code, bulk storage, and cache memory which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

5

Input/Output or I/O devices (including, but not limited to, keyboards, displays, pointing devices, DASD, tape, CDs, DVDs, thumb drives and other memory media, etc.) can be coupled to the system either directly or through intervening I/O controllers. Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modems, and Ethernet cards are just a few of the available types of network adapters.

10

15

20

25

30

Referring to FIG. 15, representative components of a Host Computer system 5000 to implement one or more aspects of the present invention are portrayed. The representative host computer 5000 comprises one or more CPUs 5001 in communication with computer memory (i.e., central storage) 5002, as well as I/O interfaces to storage media devices 5011 and networks 5010 for communicating with other computers or SANs and the like. The CPU 5001 is compliant with an architecture having an architected instruction set and architected functionality. The CPU 5001 may have dynamic address translation (DAT) 5003 for transforming program addresses (virtual addresses) into real addresses of memory. A DAT typically includes a translation lookaside buffer (TLB) 5007 for caching translations so that later accesses to the block of computer memory 5002 do not require the delay of address translation. Typically, a cache 5009 is employed between computer memory 5002 and the processor 5001. The cache 5009 may be hierarchical having a large cache available to more than one CPU and smaller, faster (lower level) caches between the large cache and each CPU. In some implementations, the lower level caches are split to provide separate low level caches for instruction fetching and data accesses. In one embodiment, an instruction is fetched from memory 5002 by an instruction fetch unit 5004 via a cache 5009. The instruction is decoded in an instruction decode unit 5006 and dispatched (with other instructions in some embodiments) to instruction execution unit or units 5008. Typically several execution units 5008 are employed, for example an arithmetic execution unit, a

floating point execution unit and a branch instruction execution unit. The instruction is executed by the execution unit, accessing operands from instruction specified registers or memory as needed. If an operand is to be accessed (loaded or stored) from memory 5002, a load/store unit 5005 typically handles the access under control of the instruction being  
5 executed. Instructions may be executed in hardware circuits or in internal microcode (firmware) or by a combination of both.

As noted, a computer system includes information in local (or main) storage, as well as addressing, protection, and reference and change recording. Some aspects of addressing  
10 include the format of addresses, the concept of address spaces, the various types of addresses, and the manner in which one type of address is translated to another type of address. Some of main storage includes permanently assigned storage locations. Main storage provides the system with directly addressable fast-access storage of data. Both data and programs are to be loaded into main storage (from input devices) before they can be  
15 processed.

Main storage may include one or more smaller, faster-access buffer storages, sometimes called caches. A cache is typically physically associated with a CPU or an I/O processor. The effects, except on performance, of the physical construction and use of distinct storage  
20 media are generally not observable by the program.

Separate caches may be maintained for instructions and for data operands. Information within a cache is maintained in contiguous bytes on an integral boundary called a cache block or cache line (or line, for short). A model may provide an EXTRACT CACHE  
25 ATTRIBUTE instruction which returns the size of a cache line in bytes. A model may also provide PREFETCH DATA and PREFETCH DATA RELATIVE LONG instructions which effects the prefetching of storage into the data or instruction cache or the releasing of data from the cache.

Storage is viewed as a long horizontal string of bits. For most operations, accesses to storage proceed in a left-to-right sequence. The string of bits is subdivided into units of eight bits. An eight-bit unit is called a byte, which is the basic building block of all information  
30

formats. Each byte location in storage is identified by a unique nonnegative integer, which is the address of that byte location or, simply, the byte address. Adjacent byte locations have consecutive addresses, starting with 0 on the left and proceeding in a left-to-right sequence. Addresses are unsigned binary integers and are 24, 31, or 64 bits.

5

Information is transmitted between storage and a CPU or a channel subsystem one byte, or a group of bytes, at a time. Unless otherwise specified, in, for instance, the z/Architecture®, a group of bytes in storage is addressed by the leftmost byte of the group. The number of bytes in the group is either implied or explicitly specified by the operation to be performed.

10

When used in a CPU operation, a group of bytes is called a field. Within each group of bytes, in, for instance, the z/Architecture®, bits are numbered in a left-to-right sequence. In the z/Architecture®, the leftmost bits are sometimes referred to as the “high-order” bits and the rightmost bits as the “low-order” bits. Bit numbers are not storage addresses, however.

15

Only bytes can be addressed. To operate on individual bits of a byte in storage, the entire byte is accessed. The bits in a byte are numbered 0 through 7, from left to right (in, e.g., the z/Architecture®). The bits in an address may be numbered 8-31 or 40-63 for 24-bit addresses, or 1-31 or 33-63 for 31-bit addresses; they are numbered 0-63 for 64-bit

20

addresses. Within any other fixed-length format of multiple bytes, the bits making up the format are consecutively numbered starting from 0. For purposes of error detection, and in preferably for correction, one or more check bits may be transmitted with each byte or with a group of bytes. Such check bits are generated automatically by the machine and cannot be directly controlled by the program. Storage capacities are expressed in number of bytes.

25

When the length of a storage-operand field is implied by the operation code of an instruction, the field is said to have a fixed length, which can be one, two, four, eight, or sixteen bytes. Larger fields may be implied for some instructions. When the length of a storage-operand field is not implied but is stated explicitly, the field is said to have a variable length. Variable-length operands can vary in length by increments of one byte (or with some instructions, in multiples of two bytes or other multiples). When information is placed in storage, the contents of only those byte locations are replaced that are included in the designated field, even though the width of the physical path to storage may be greater than the length of the field being stored.

30

Certain units of information are to be on an integral boundary in storage. A boundary is called integral for a unit of information when its storage address is a multiple of the length of the unit in bytes. Special names are given to fields of 2, 4, 8, and 16 bytes on an integral boundary. A halfword is a group of two consecutive bytes on a two-byte boundary and is the basic building block of instructions. A word is a group of four consecutive bytes on a four-byte boundary. A doubleword is a group of eight consecutive bytes on an eight-byte boundary. A quadword is a group of 16 consecutive bytes on a 16-byte boundary. When storage addresses designate halfwords, words, doublewords, and quadwords, the binary representation of the address contains one, two, three, or four rightmost zero bits, respectively. Instructions are to be on two-byte integral boundaries. The storage operands of most instructions do not have boundary-alignment requirements.

On devices that implement separate caches for instructions and data operands, a significant delay may be experienced if the program stores into a cache line from which instructions are subsequently fetched, regardless of whether the store alters the instructions that are subsequently fetched.

In one embodiment, the invention may be practiced by software (sometimes referred to licensed internal code, firmware, micro-code, milli-code, pico-code and the like, any of which would be consistent with the present invention). Referring to FIG. 15, software program code which embodies the present invention is typically accessed by processor 5001 of the host system 5000 from long-term storage media devices 5011, such as a CD-ROM drive, tape drive or hard drive. The software program code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from computer memory 5002 or storage of one computer system over a network 5010 to other computer systems for use by users of such other systems.

The software program code includes an operating system which controls the function and interaction of the various computer components and one or more application programs. Program code is normally paged from storage media device 5011 to the relatively higher-speed computer storage 5002 where it is available for processing by processor 5001. The

techniques and methods for embodying software program code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein. Program code, when created and stored on a tangible medium (including but not limited to electronic memory modules (RAM), flash memory, Compact Discs (CDs),  
5 DVDs, Magnetic Tape and the like is often referred to as a “computer program product”. The computer program product medium is typically readable by a processing circuit preferably in a computer system for execution by the processing circuit.

FIG. 16 illustrates a representative workstation or server hardware system in which the  
10 present invention may be practiced. The system 5020 of FIG. 16 comprises a representative base computer system 5021, such as a personal computer, a workstation or a server, including optional peripheral devices. The base computer system 5021 includes one or more processors 5026 and a bus employed to connect and enable communication between the processor(s) 5026 and the other components of the system 5021 in accordance with known  
15 techniques. The bus connects the processor 5026 to memory 5025 and long-term storage 5027 which can include a hard drive (including any of magnetic media, CD, DVD and Flash Memory for example) or a tape drive for example. The system 5021 might also include a user interface adapter, which connects the microprocessor 5026 via the bus to one or more interface devices, such as a keyboard 5024, a mouse 5023, a printer/scanner 5030 and/or  
20 other interface devices, which can be any user interface device, such as a touch sensitive screen, digitized entry pad, etc. The bus also connects a display device 5022, such as an LCD screen or monitor, to the microprocessor 5026 via a display adapter.

The system 5021 may communicate with other computers or networks of computers by way  
25 of a network adapter capable of communicating 5028 with a network 5029. Example network adapters are communications channels, token ring, Ethernet or modems. Alternatively, the system 5021 may communicate using a wireless interface, such as a CDPD (cellular digital packet data) card. The system 5021 may be associated with such other computers in a Local Area Network (LAN) or a Wide Area Network (WAN), or the system  
30 5021 can be a client in a client/server arrangement with another computer, etc. All of these configurations, as well as the appropriate communications hardware and software, are known in the art.

FIG. 17 illustrates a data processing network 5040 in which the present invention may be practiced. The data processing network 5040 may include a plurality of individual networks, such as a wireless network and a wired network, each of which may include a plurality of individual workstations 5041, 5042, 5043, 5044. Additionally, as those skilled in the art will appreciate, one or more LANs may be included, where a LAN may comprise a plurality of intelligent workstations coupled to a host processor.

Still referring to FIG. 17, the networks may also include mainframe computers or servers, such as a gateway computer (client server 5046) or application server (remote server 5048 which may access a data repository and may also be accessed directly from a workstation 5045). A gateway computer 5046 serves as a point of entry into each individual network. A gateway is needed when connecting one networking protocol to another. The gateway 5046 may be preferably coupled to another network (the Internet 5047 for example) by means of a communications link. The gateway 5046 may also be directly coupled to one or more workstations 5041, 5042, 5043, 5044 using a communications link. The gateway computer may be implemented utilizing an IBM eServer™ System z® server available from International Business Machines Corporation.

Referring concurrently to FIG. 16 and FIG. 17, software programming code which may embody the present invention may be accessed by the processor 5026 of the system 5020 from long-term storage media 5027, such as a CD-ROM drive or hard drive. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users 5050, 5051 from the memory or storage of one computer system over a network to other computer systems for use by users of such other systems.

Alternatively, the programming code may be embodied in the memory 5025, and accessed by the processor 5026 using the processor bus. Such programming code includes an operating system which controls the function and interaction of the various computer components and one or more application programs 5032. Program code is normally paged from storage media 5027 to high-speed memory 5025 where it is available for processing by

the processor 5026. The techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein. Program code, when created and stored on a tangible medium (including but not limited to electronic memory modules (RAM), flash memory, Compact Discs (CDs), DVDs, Magnetic Tape and the like is often referred to as a “computer program product”. The computer program product medium is typically readable by a processing circuit preferably in a computer system for execution by the processing circuit.

The cache that is most readily available to the processor (normally faster and smaller than other caches of the processor) is the lowest (L1 or level one) cache and main store (main memory) is the highest level cache (L3 if there are 3 levels). The lowest level cache is often divided into an instruction cache (I-Cache) holding machine instructions to be executed and a data cache (D-Cache) holding data operands.

Referring to FIG. 18, an exemplary processor embodiment is depicted for processor 5026. Typically one or more levels of cache 5053 are employed to buffer memory blocks in order to improve processor performance. The cache 5053 is a high speed buffer holding cache lines of memory data that are likely to be used. Typical cache lines are 64, 128 or 256 bytes of memory data. Separate caches are often employed for caching instructions than for caching data. Cache coherence (synchronization of copies of lines in memory and the caches) is often provided by various “snoop” algorithms well known in the art. Main memory storage 5025 of a processor system is often referred to as a cache. In a processor system having 4 levels of cache 5053, main storage 5025 is sometimes referred to as the level 5 (L5) cache since it is typically faster and only holds a portion of the non-volatile storage (DASD, tape etc) that is available to a computer system. Main storage 5025 “caches” pages of data paged in and out of the main storage 5025 by the operating system.

A program counter (instruction counter) 5061 keeps track of the address of the current instruction to be executed. A program counter in a z/Architecture<sup>®</sup> processor is 64 bits and can be truncated to 31 or 24 bits to support prior addressing limits. A program counter is typically embodied in a PSW (program status word) of a computer such that it persists

during context switching. Thus, a program in progress, having a program counter value, may be interrupted by, for example, the operating system (context switch from the program environment to the operating system environment). The PSW of the program maintains the program counter value while the program is not active, and the program counter (in the PSW) of the operating system is used while the operating system is executing. Typically, the program counter is incremented by an amount equal to the number of bytes of the current instruction. RISC (Reduced Instruction Set Computing) instructions are typically fixed length while CISC (Complex Instruction Set Computing) instructions are typically variable length. Instructions of the IBM z/Architecture<sup>®</sup> are CISC instructions having a length of 2, 4 or 6 bytes. The Program counter 5061 is modified by either a context switch operation or a branch taken operation of a branch instruction for example. In a context switch operation, the current program counter value is saved in the program status word along with other state information about the program being executed (such as condition codes), and a new program counter value is loaded pointing to an instruction of a new program module to be executed. A branch taken operation is performed in order to permit the program to make decisions or loop within the program by loading the result of the branch instruction into the program counter 5061.

Typically an instruction fetch unit 5055 is employed to fetch instructions on behalf of the processor 5026. The fetch unit either fetches "next sequential instructions", target instructions of branch taken instructions, or first instructions of a program following a context switch. Modern Instruction fetch units often employ prefetch techniques to speculatively prefetch instructions based on the likelihood that the prefetched instructions might be used. For example, a fetch unit may fetch 16 bytes of instruction that includes the next sequential instruction and additional bytes of further sequential instructions.

The fetched instructions are then executed by the processor 5026. In an embodiment, the fetched instruction(s) are passed to a dispatch unit 5056 of the fetch unit. The dispatch unit decodes the instruction(s) and forwards information about the decoded instruction(s) to appropriate units 5057, 5058, 5060. An execution unit 5057 will typically receive information about decoded arithmetic instructions from the instruction fetch unit 5055 and will perform arithmetic operations on operands according to the opcode of the instruction.

Operands are provided to the execution unit 5057 preferably either from memory 5025, architected registers 5059 or from an immediate field of the instruction being executed. Results of the execution, when stored, are stored either in memory 5025, registers 5059 or in other machine hardware (such as control registers, PSW registers and the like).

5

A processor 5026 typically has one or more units 5057, 5058, 5060 for executing the function of the instruction. Referring to FIG. 19A, an execution unit 5057 may communicate with architected general registers 5059, a decode/dispatch unit 5056, a load store unit 5060, and other 5065 processor units by way of interfacing logic 5071. An execution unit 5057 may employ several register circuits 5067, 5068, 5069 to hold information that the arithmetic logic unit (ALU) 5066 will operate on. The ALU performs arithmetic operations such as add, subtract, multiply and divide as well as logical function such as and, or and exclusive-or (XOR), rotate and shift. Preferably the ALU supports specialized operations that are design dependent. Other circuits may provide other architected facilities 5072 including condition codes and recovery support logic for example. Typically the result of an ALU operation is held in an output register circuit 5070 which can forward the result to a variety of other processing functions. There are many arrangements of processor units, the present description is only intended to provide a representative understanding of one embodiment.

20

An ADD instruction for example would be executed in an execution unit 5057 having arithmetic and logical functionality while a floating point instruction for example would be executed in a floating point execution having specialized floating point capability. Preferably, an execution unit operates on operands identified by an instruction by performing an opcode defined function on the operands. For example, an ADD instruction may be executed by an execution unit 5057 on operands found in two registers 5059 identified by register fields of the instruction.

25

The execution unit 5057 performs the arithmetic addition on two operands and stores the result in a third operand where the third operand may be a third register or one of the two source registers. The execution unit preferably utilizes an Arithmetic Logic Unit (ALU) 5066 that is capable of performing a variety of logical functions such as Shift, Rotate, And,

30

Or and XOR as well as a variety of algebraic functions including any of add, subtract, multiply, divide. Some ALUs 5066 are designed for scalar operations and some for floating point. Data may be Big Endian (where the least significant byte is at the highest byte address) or Little Endian (where the least significant byte is at the lowest byte address) depending on architecture. The IBM z/Architecture<sup>®</sup> is Big Endian. Signed fields may be sign and magnitude, 1's complement or 2's complement depending on architecture. A 2's complement number is advantageous in that the ALU does not need to design a subtract capability since either a negative value or a positive value in 2's complement requires only an addition within the ALU. Numbers are commonly described in shorthand, where a 12 bit field defines an address of a 4,096 byte block and is commonly described as a 4 Kbyte (Kilo-byte) block, for example.

Referring to FIG. 19B, branch instruction information for executing a branch instruction is typically sent to a branch unit 5058 which often employs a branch prediction algorithm such as a branch history table 5082 to predict the outcome of the branch before other conditional operations are complete. The target of the current branch instruction will be fetched and speculatively executed before the conditional operations are complete. When the conditional operations are completed the speculatively executed branch instructions are either completed or discarded based on the conditions of the conditional operation and the speculated outcome. A typical branch instruction may test condition codes and branch to a target address if the condition codes meet the branch requirement of the branch instruction, a target address may be calculated based on several numbers including ones found in register fields or an immediate field of the instruction for example. The branch unit 5058 may employ an ALU 5074 having a plurality of input register circuits 5075, 5076, 5077 and an output register circuit 5080. The branch unit 5058 may communicate with general registers 5059, decode dispatch unit 5056 or other circuits 5073, for example.

The execution of a group of instructions can be interrupted for a variety of reasons including a context switch initiated by an operating system, a program exception or error causing a context switch, an I/O interruption signal causing a context switch or multi-threading activity of a plurality of programs (in a multi-threaded environment), for example. Preferably a context switch action saves state information about a currently executing program and then

loads state information about another program being invoked. State information may be saved in hardware registers or in memory for example. State information preferably comprises a program counter value pointing to a next instruction to be executed, condition codes, memory translation information and architected register content. A context switch activity can be exercised by hardware circuits, application programs, operating system programs or firmware code (microcode, pico-code or licensed internal code (LIC)) alone or in combination.

A processor accesses operands according to instruction defined methods. The instruction may provide an immediate operand using the value of a portion of the instruction, may provide one or more register fields explicitly pointing to either general purpose registers or special purpose registers (floating point registers for example). The instruction may utilize implied registers identified by an opcode field as operands. The instruction may utilize memory locations for operands. A memory location of an operand may be provided by a register, an immediate field, or a combination of registers and immediate field as exemplified by the z/Architecture<sup>®</sup> long displacement facility wherein the instruction defines a base register, an index register and an immediate field (displacement field) that are added together to provide the address of the operand in memory for example. Location herein typically implies a location in main memory (main storage) unless otherwise indicated.

Referring to FIG. 19C, a processor accesses storage using a load/store unit 5060. The load/store unit 5060 may perform a load operation by obtaining the address of the target operand in memory 5053 and loading the operand in a register 5059 or another memory 5053 location, or may perform a store operation by obtaining the address of the target operand in memory 5053 and storing data obtained from a register 5059 or another memory 5053 location in the target operand location in memory 5053. The load/store unit 5060 may be speculative and may access memory in a sequence that is out-of-order relative to instruction sequence, however the load/store unit 5060 is to maintain the appearance to programs that instructions were executed in order. A load/store unit 5060 may communicate with general registers 5059, decode/dispatch unit 5056, cache/memory interface 5053 or other elements 5083 and comprises various register circuits, ALUs 5085 and control logic 5090 to calculate storage addresses and to provide pipeline sequencing to keep operations in-

order. Some operations may be out of order but the load/store unit provides functionality to make the out of order operations to appear to the program as having been performed in order, as is well known in the art.

5 Preferably addresses that an application program “sees” are often referred to as virtual addresses. Virtual addresses are sometimes referred to as “logical addresses” and “effective addresses”. These virtual addresses are virtual in that they are redirected to physical  
10 memory location by one of a variety of dynamic address translation (DAT) technologies including, but not limited to, simply prefixing a virtual address with an offset value, translating the virtual address via one or more translation tables, the translation tables preferably comprising at least a segment table and a page table alone or in combination, preferably, the segment table having an entry pointing to the page table. In the  
15 z/Architecture<sup>®</sup>, a hierarchy of translation is provided including a region first table, a region second table, a region third table, a segment table and an optional page table. The performance of the address translation is often improved by utilizing a translation lookaside buffer (TLB) which comprises entries mapping a virtual address to an associated physical  
20 memory location. The entries are created when the DAT translates a virtual address using the translation tables. Subsequent use of the virtual address can then utilize the entry of the fast TLB rather than the slow sequential translation table accesses. TLB content may be managed by a variety of replacement algorithms including LRU (Least Recently used).

In the case where the processor is a processor of a multi-processor system, each processor has responsibility to keep shared resources, such as I/O, caches, TLBs and memory,  
25 interlocked for coherency. Typically, “snoop” technologies will be utilized in maintaining cache coherency. In a snoop environment, each cache line may be marked as being in any one of a shared state, an exclusive state, a changed state, an invalid state and the like in order to facilitate sharing.

I/O units 5054 (FIG. 18) provide the processor with means for attaching to peripheral  
30 devices including tape, disc, printers, displays, and networks for example. I/O units are often presented to the computer program by software drivers. In mainframes, such as the System

z<sup>®</sup> from IBM<sup>®</sup>, channel adapters and open system adapters are I/O units of the mainframe that provide the communications between the operating system and peripheral devices.

Further, other types of computing environments can benefit from one or more aspects of the present invention. As an example, an environment may include an emulator (e.g., software or other emulation mechanisms), in which a particular architecture (including, for instance, instruction execution, architected functions, such as address translation, and architected registers) or a subset thereof is emulated (e.g., on a native computer system having a processor and memory). In such an environment, one or more emulation functions of the emulator can implement one or more aspects of the present invention, even though a computer executing the emulator may have a different architecture than the capabilities being emulated. As one example, in emulation mode, the specific instruction or operation being emulated is decoded, and an appropriate emulation function is built to implement the individual instruction or operation.

In an emulation environment, a host computer includes, for instance, a memory to store instructions and data; an instruction fetch unit to fetch instructions from memory and to optionally, provide local buffering for the fetched instruction; an instruction decode unit to receive the fetched instructions and to determine the type of instructions that have been fetched; and an instruction execution unit to execute the instructions. Execution may include loading data into a register from memory; storing data back to memory from a register; or performing some type of arithmetic or logical operation, as determined by the decode unit. In one example, each unit is implemented in software. For instance, the operations being performed by the units are implemented as one or more subroutines within emulator software.

More particularly, in a mainframe, architected machine instructions are used by programmers, usually today "C" programmers, often by way of a compiler application. These instructions stored in the storage medium may be executed natively in a z/Architecture<sup>®</sup> IBM<sup>®</sup> Server, or alternatively in machines executing other architectures. They can be emulated in the existing and in future IBM<sup>®</sup> mainframe servers and on other machines of IBM<sup>®</sup> (e.g., Power Systems servers and System x<sup>®</sup> Servers). They can be

executed in machines running Linux on a wide variety of machines using hardware manufactured by IBM<sup>®</sup>, Intel<sup>®</sup>, AMD<sup>™</sup>, and others. Besides execution on that hardware under a z/Architecture<sup>®</sup>, Linux can be used as well as machines which use emulation by Hercules (see [www.hercules-390.org](http://www.hercules-390.org)), or FSI (Fundamental Software, Inc) (see [www.funsoft.com](http://www.funsoft.com)), where generally execution is in an emulation mode. In emulation mode, emulation software is executed by a native processor to emulate the architecture of an emulated processor.

The native processor typically executes emulation software comprising either firmware or a native operating system to perform emulation of the emulated processor. The emulation software is responsible for fetching and executing instructions of the emulated processor architecture. The emulation software maintains an emulated program counter to keep track of instruction boundaries. The emulation software may fetch one or more emulated machine instructions at a time and convert the one or more emulated machine instructions to a corresponding group of native machine instructions for execution by the native processor. These converted instructions may be cached such that a faster conversion can be accomplished. Notwithstanding, the emulation software is to maintain the architecture rules of the emulated processor architecture so as to assure operating systems and applications written for the emulated processor operate correctly. Furthermore, the emulation software is to provide resources identified by the emulated processor architecture including, but not limited to, control registers, general purpose registers, floating point registers, dynamic address translation function including segment tables and page tables for example, interrupt mechanisms, context switch mechanisms, Time of Day (TOD) clocks and architected interfaces to I/O subsystems such that an operating system or an application program designed to run on the emulated processor, can be run on the native processor having the emulation software.

A specific instruction being emulated is decoded, and a subroutine is called to perform the function of the individual instruction. An emulation software function emulating a function of an emulated processor is implemented, for example, in a "C" subroutine or driver, or some other method of providing a driver for the specific hardware as will be within the skill of those in the art after understanding the description of the preferred embodiment. Various

software and hardware emulation patents including, but not limited to U.S. Letters Patent No. 5,551,013, entitled "Multiprocessor for Hardware Emulation", by Beausoleil et al.; and U.S. Letters Patent No. 6,009,261, entitled "Preprocessing of Stored Target Routines for Emulating Incompatible Instructions on a Target Processor", by Scalzi et al; and U.S. Letters Patent No. 5,574,873, entitled "Decoding Guest Instruction to Directly Access Emulation Routines that Emulate the Guest Instructions", by Davidian et al; and U.S. Letters Patent No. 6,308,255, entitled "Symmetrical Multiprocessing Bus and Chipset Used for Coprocessor Support Allowing Non-Native Code to Run in a System", by Gorishek et al; and U.S. Letters Patent No. 6,463,582, entitled "Dynamic Optimizing Object Code Translator for Architecture Emulation and Dynamic Optimizing Object Code Translation Method", by Lethin et al; and U.S. Letters Patent No. 5,790,825, entitled "Method for Emulating Guest Instructions on a Host Computer Through Dynamic Recompilation of Host Instructions", by Eric Traut; and many others, illustrate a variety of known ways to achieve emulation of an instruction format architected for a different machine for a target machine available to those skilled in the art.

In FIG. 20, an example of an emulated host computer system 5092 is provided that emulates a host computer system 5000' of a host architecture. In the emulated host computer system 5092, the host processor (CPU) 5091 is an emulated host processor (or virtual host processor) and comprises an emulation processor 5093 having a different native instruction set architecture than that of the processor 5091 of the host computer 5000'. The emulated host computer system 5092 has memory 5094 accessible to the emulation processor 5093. In the example embodiment, the memory 5094 is partitioned into a host computer memory 5096 portion and an emulation routines 5097 portion. The host computer memory 5096 is available to programs of the emulated host computer 5092 according to host computer architecture. The emulation processor 5093 executes native instructions of an architected instruction set of an architecture other than that of the emulated processor 5091, the native instructions obtained from emulation routines memory 5097, and may access a host instruction for execution from a program in host computer memory 5096 by employing one or more instruction(s) obtained in a sequence & access/decode routine which may decode the host instruction(s) accessed to determine a native instruction execution routine for emulating the function of the host instruction accessed. Other facilities that are defined for the host

computer system 5000' architecture may be emulated by architected facilities routines, including such facilities as general purpose registers, control registers, dynamic address translation and I/O subsystem support and processor cache, for example. The emulation routines may also take advantage of functions available in the emulation processor 5093 (such as general registers and dynamic translation of virtual addresses) to improve performance of the emulation routines. Special hardware and off-load engines may also be provided to assist the processor 5093 in emulating the function of the host computer 5000'.

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising", when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below, if any, are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the principles of the invention and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiment with various modifications as are suited to the particular use contemplated.

## CLAIMS

1. A computer program product for managing interruption requests in a computing environment, the computer program product comprising:
  - a non-transitory computer readable storage medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising:
    - based on executing a register interruptions operation of a Modify PCI Function Controls (MPFC) instruction that specifies a function handle of an adapter, specifying a location in system memory of an adapter interruption bit vector (AIBV) of the adapter, the AIBV included in an array of one or more AIBVs, and a location in system memory of an adapter interruption summary bit (AISB) of an AISB array;
    - receiving from the adapter a request for interruption; and
    - based on the received request, setting, by an input/output (I/O) hub coupled to the adapter, an indicator in the AIBV indicating a type of event from the adapter and setting the AISB indicating an indicator is set in the AIBV, wherein the setting the indicator in the AIBV comprises: determining whether a vector number provided in the request is within a number of interruptions allowed for the adapter; and based on determining the vector number is within the allowed number of interruptions, using the vector number, an AIBV offset, and an AIBV address to determine the indicator in the AIBV to be set.
2. The computer program product of claim 1, wherein the method further comprises presenting an interruption to an operating system, the interruption based on the request for interruption.
3. The computer program product of claim 2, wherein the request for interruption represents a plurality of message signaled interruptions, and the interruption to the operating system is part of an input/output adapter event notification for the operating system.
4. The computer program product of claim 2, wherein the method further comprises obtaining, based on the presenting, one or more AIBV indications for one or more adapters specifying at least one reason for interruption per adapter.

5. The computer program product of claim 4, wherein the obtaining comprises obtaining a plurality of AIBV indications specifying a plurality of reasons for interruption, the plurality of reasons for interruption corresponding to the plurality of requests for interruption.
6. The computer program product of claim 5, wherein the obtaining comprises using the AISB in obtaining the one or more AIBV indications in the one or more AIBVs.
7. The computer program product of claim 1, wherein the setting the indicator in the AIBV comprises: using an identifier of the request from the adapter to obtain a device table entry, the device table entry comprising a value specifying the number of interruptions that are allowed for the adapter; and based on determining the vector number is within the allowed number of interruptions, using one or more parameters of the device table entry to locate a starting position of the AIBV.
8. A computer system for managing interruption requests in a computing environment, the computer system comprising:
  - a memory; and
  - a processor in communications with the memory, wherein the computer system is configured to perform a method, the method comprising:
    - based on executing a register interruptions operation of a Modify PCI Function Controls (MPFC) instruction that specifies a function handle of an adapter, specifying a location in system memory of an adapter interruption bit vector (AIBV) of the adapter, the AIBV included in an array of one or more AIBVs, and a location in system memory of an adapter interruption summary bit (AISB) of an AISB array;
    - receiving from the adapter a request for interruption; and
    - based on the received request, setting, by an input/output (I/O) hub coupled to the adapter, an indicator in the AIBV indicating a type of event from the adapter and setting the AISB indicating an indicator is set in the AIBV, wherein the setting the indicator in the AIBV comprises:
      - determining whether a vector number provided in the request is within a number of interruptions allowed for the adapter; and

based on determining the vector number is within the allowed number of interruptions, using the vector number, an AIBV offset, and an AIBV address to determine the indicator in the AIBV to be set.

9. The computer system of claim 8, wherein the method further comprises presenting an interruption to an operating system, the interruption based on the request for interruption.
10. The computer system of claim 9, wherein the request for interruption represents a plurality of message signaled interruptions, and the interruption to the operating system is part of an input/output adapter event notification for the operating system.
11. The computer system of claim 9, wherein the method further comprises obtaining, based on the presenting, one or more AIBV indications for one or more adapters specifying at least one reason for interruption per adapter.
12. The computer system of claim 11, wherein the obtaining comprises obtaining a plurality of AIBV indications specifying a plurality of reasons for interruption, the plurality of reasons for interruption corresponding to the plurality of requests for interruption.
13. The computer system of claim 12, wherein the obtaining comprises using the AISB in obtaining the one or more AIBV indications in the one or more AIBVs.
14. The computer system of claim 8, wherein the setting the indicator in the AIBV comprises: using an identifier of the request from the adapter to obtain a device table entry, the device table entry comprising a value specifying the number of interruptions that are allowed for the adapter; and based on determining the vector number is within the allowed number of interruptions, using one or more parameters of the device table entry to locate a starting position of the AIBV.
15. The computer program product of claim 1, wherein the AIBV for the adapter comprises a plurality of indicators for the adapter, the plurality of indicators indicating a plurality of types of events associated with the adapter.
16. The computer program product of claim 1, wherein the array of one or more AIBVs includes another AIBV for another adapter, and wherein a number of indicators

in the AIBV for the adapter is different from the number of indicators in the another AIBV for the another adapter.

17. The computer program product of claim 1, wherein the function handle includes an enable indicator that indicates whether the handle is enabled, a number that identifies the adapter, and an instance number to indicate a particular instance of the handle.

18. The computer system of claim 8, wherein the AIBV for the adapter comprises a plurality of indicators for the adapter, the plurality of indicators indicating a plurality of types of events associated with the adapter.

19. The computer system of claim 8, wherein the array of one or more AIBVs includes another AIBV for another adapter, and wherein a number of indicators in the AIBV for the adapter is different from the number of indicators in the another AIBV for the another adapter.

20. The computer system of claim 8, wherein the function handle includes an enable indicator that indicates whether the handle is enabled, a number that identifies the adapter, and an instance number to indicate a particular instance of the handle.