

White Paper

ABSTRACT

This white paper presents a description of the Magic Packet Technology and how it works. It also covers some issues involving the sleeping Green PC and how the Magic Packet Technology can be used to put a PC in a low-power state and still be manageable by a network system administrator.

SCOPE

The PC market has many forces influencing the design and implementation of PCs, operating systems, and peripheral devices. Most of the time, these forces are complementary, such as CPU performance and multimedia uses, or the desire for end users to have their machines backed up each night, and the Information Systems (IS) department's desire to maintain data integrity throughout the network.

However, sometimes the forces working on the PC market seem to be in direct conflict. One such example is the desire of the IS department to be able to do end node management (software updates, backups, etc.) and the US Government's desire, through the Energy Star Program, to put PCs to sleep soon after the PC is not in use. If the PC falls asleep at 5:30 in the evening, how will the IS department manage it? If the user feels it is more important to have his hard disk backed up than save power, will he disable the power savings feature of his PC? If so, this action will nullify the Energy Star Program's intent to save power by PCs putting themselves to sleep when not in use.

PROPOSAL

AMD and Hewlett Packard identified this problem almost 2 years ago and began working to find a solution to the problem of the networked Green PC (*Green* being the industry term for a PC that goes into low-power mode upon sensing no activity). This collaboration resulted in a proposal for an industry standard mechanism that allows a networked PC to go completely asleep (Deep Green), yet still allows the network administrator or some kind of network management software to wake up the PC simply by sending it a specific Ethernet frame. This standard Ethernet frame contains a specific data pattern detected by the Ethernet controller on the receiving end. The Ethernet controller then alerts the system and the power management circuitry wakes it up.

Below are details on the silicon implementation of Magic Packet Technology, as well as some of the system and software implications.

Magic Packet Technology Overview

The basic technical details of Magic Packet Technology are simple and easy to understand. There is also a second set of details, which will be implementation specific. In other words, silicon- or gate-level implementations of Magic Packet Technology may differ from AMD's approach and be completely interoperable, as long as the basic feature set is maintained.

Magic Packet Technology is a feature designed to be incorporated into an Ethernet controller. The basic feature set consists of the following:

- Magic Packet Mode Enable
- Magic Packet Frame Detection
- Magic Packet Mode Disable

Magic Packet Mode Enable

Assuming the Ethernet controller is running and communicating with the network, there must be a way for the PC's power management hardware or software to put the Ethernet controller into the Magic Packet mode prior to the system going to sleep. On both the PCnet-ISA II and PCnet-PCI II devices, this can be accomplished two ways: either by setting a bit in an internal register or by driving the SLEEP# pin low. Either one of these actions will disable normal network activity and enable Magic Packet mode. The device will no longer generate any transmits and will monitor all incoming frames to determine if any of them is a Magic Packet frame.

Magic Packet Frame Detection

Once the LAN controller has been put into the Magic Packet mode, it scans all incoming frames addressed to the node for a specific data sequence, which indicates to the controller that this is a Magic Packet frame. A Magic Packet frame must also meet the basic requirements for the LAN technology chosen, such as

SOURCE ADDRESS, DESTINATION ADDRESS (which may be the receiving station's IEEE address or a MULTICAST address which includes the BROADCAST address), and CRC. The specific sequence consists of 16 duplications of the IEEE address of this node, with no breaks or interruptions.

This sequence can be located anywhere within the packet, but must be preceded by a synchronization stream. The synchronization stream allows the scanning state machine to be much simpler. The synchronization stream is defined as 6 bytes of FFh. The device will also accept a MULTICAST frame, as long as the 16 duplications of the IEEE address match the address of the machine to be awakened.

If the IEEE address for a particular node on the network was 11h 22h 33h 44h 55h 66h, then the LAN controller would be scanning for the data sequence (assuming an Ethernet Frame):

```

DESTINATION SOURCE MISC FF FF FF FF FF
FF 11 22 33 44 55 66 11 22 33 44 55 66 11 22 33 44
55 66 11 22 33 44 55 66 11 22 33 44 55 66 11 22 33
44 55 66 11 22 33 44 55 66 11 22 33 44 55 66 11 22
33 44 55 66 11 22 33 44 55 66 11 22 33 44 55 66 11
22 33 44 55 66 11 22 33 44 55 66 11 22 33 44 55 66
11 22 33 44 55 66 11 22 33 44 55 66 MISC CRC
    
```

Magic Packet Mode Disable

There are two instances where the Magic Packet mode must be disabled, and the Ethernet controller returned to normal operation mode. Either the system has received a Magic Packet frame, possibly from a network administrator who wants to do a hard disk backup, or some other action has caused the system to leave the low power sleep state, such as a user touching the keyboard, moving the mouse, etc. In either case, the power management hardware or software must disable the Magic Packet mode and return the Ethernet controller to normal operation. On the PCnet-ISA II or the PCnet-PCI II, this may be accomplished by either resetting the register bit in an internal register, or de-asserting the SLEEP# pin.

SILICON IMPLEMENTATION

Implementation details of the Magic Packet Technology may vary from device to device, as long as the basic functionality is maintained. Since an Ethernet controller already has built-in address matching circuitry to recognize regular frames addressed to the node, this circuitry may be re-used in the case of Magic Packet mode. A new mode of operation must be implemented, which will allow the power management software or hardware to enter and leave Magic Packet mode. A counter must be added to the address matching circuitry to count up the 16 duplications of the IEEE address, with another circuit to reset the counter if the data being processed does not match the IEEE address.

It should not take much design effort, time, or silicon area to add the Magic Packet Technology to an existing Ethernet controller. This was one of the primary reasons for going with a solution that uses the IEEE address as the identifier for the Magic Packet frame, since the circuitry already exists to match this data stream.

SYSTEM IMPLEMENTATION

To utilize the Magic Packet Technology in a PC, there are several modifications which must be done to ensure proper operation of the feature. Let's take the case of a motherboard implementation, where the Ethernet controller will be located on the motherboard, with an RJ-45 connector coming out the back of the computer. This implementation is becoming more and more common as Ethernet establishes itself as a de facto standard for the desktop LAN.

First, let's address the hardware steps necessary to allow the Ethernet controller to wake up the system when it receives a Magic Packet frame. Most desktop PC's these days already have pretty advanced power management circuitry either built into the chipset or as a separate functional block on the motherboard. In this example, it is a simple matter to connect one of the LED pins, such as LED 3, into the power management circuitry. The Magic Packet frame indication becomes just another possible alert to the power management circuitry that the system needs to wake up. The details as to how to connect the LED 3 pin into the system's power management circuitry is obviously system, chipset, and design specific, so will not be covered here in detail.

The second stage, the enabling and disabling of the Magic Packet mode, can be done in hardware or software. If done in hardware, the power management circuitry must drive the SLEEP# pin active before it places the machine in the sleep mode. This will stop all normal network activity and place the Ethernet controller into the Magic Packet mode. Upon receiving a Magic Packet frame or sensing some other activity, such as keyboard or mouse movement, the hardware must then de-assert the SLEEP# pin to remove the Ethernet controller from Magic Packet mode and return the controller to normal operation.

The second stage can also be done in software, if desired. On most systems, the BIOS or other software will be aware of the state of the system and will cooperate in the powering down of the various components of the system. In this instance, if the BIOS is involved in the powering up and down of subsystems, it could, when powering down the system to go to sleep, set a bit in the Ethernet controller to enable Magic Packet mode. When the system wakes up, for whatever reason, the BIOS would then disable Magic Packet mode by de-asserting the same bit to turn off Magic Packet frame detection feature and return the Ethernet controller to the normal operating mode.

Driver Implications

Network operating system drivers may or may not need to be modified to support the Magic Packet mode of operation. If the hardware or BIOS is responsible for enabling and disabling the Magic Packet mode, the driver may not need to be changed at all. This is because the driver has no need to know whether the system has been awake or asleep. If the system goes asleep and then wakes up for some reason, as long as the Ethernet controller is in the same state as when the driver was last accessing it, including the registers and buffer pointers, then the driver can continue functioning as if nothing happened. Of course, if the system was asleep for an extended time, the PC may have been logged off the network, but it is up to the upper layers, not the driver, to re-establish the link to the server.

However, all that said, the driver may very well be the best place to put support for Magic Packet Technology. For instance, the driver could monitor all Advanced Power Management (APM) calls, and if told that the system is going to sleep, the driver could enable Magic Packet mode. When the APM call indicates the system is in the process of waking up, the driver could then disable Magic Packet mode, verify the state of the Ethernet controller, and continue normal activity.

NOS Implications

The implications of Magic Packet Technology on the Network Operating Systems (NOS) is still being determined; however, there are several points which should be highlighted. First and foremost, many current NOS periodically send a packet to an end node to determine if the PC is still there and will log off any station that does not respond to this Ping after X number of retries. This has caused notebook users frustration for years, since notebooks are designed to go into a SUSPEND mode after a few minutes of inactivity.

This pinging and subsequent logging off of PCs that do not respond will have to change, and, in fact, is in the process of changing. The most widely used NOS which does this is Novell Interware. There are new options in Novell Netware 4.1 which will allow a network administrator to disable this function.

Another problem currently exists with Peer-to-Peer networking operating systems, such as Windows for Workgroups. The problem is that a user may attempt to log onto another user's computer after hours, when that computer has gone into Magic Packet mode and does not respond to normal network traffic. In this case, the user will not be able to attach to the other computer as a server.

Infrastructure Implications

Since implementing Magic Packet Technology, there have been many questions concerning the ability of a Magic Packet frame to actually reach and power up a

remote PC. These questions usually center around the ability of the Magic Packet frame to bridge and route to a remote PC, which may be in the next building or across the country.

Bridges

First, let's address the bridge issue. Since the Magic Packet frame is a standard Ethernet frame, there is no reason a bridge would not forward it appropriately. The only difference between a Magic Packet frame and a standard Ethernet frame is the 16 duplications of the 6-byte IEEE address, which is carried in the data portion of the frame. The bridge does not even care what is in the data portion of a frame, it only cares about the destination address. The question has also been asked "what happens if a bridge has deleted the address of the end node the Magic Packet frame is destined for from the bridge address database? Will the bridge just throw the frame away?" The answer is no. If a bridge does not know on which port a specific address is located, the bridge must forward the frame to all ports. This will guarantee that the node, which the frame has been addressed to, will receive it.

Routers

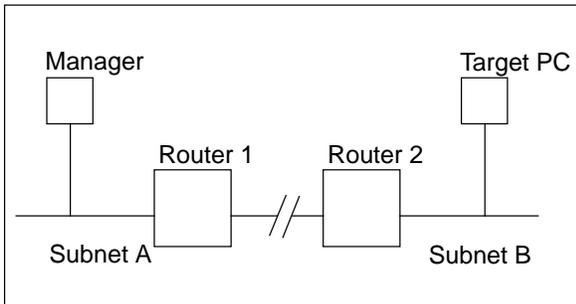
Next we will focus on the routers in a network. Many networks are separated by routers; a router is almost always used to connect a LAN in one building to a LAN in another building. Routers are used to segment the LAN and reduce the number of nodes, as well as the number of broadcast messages, on each segment of the LAN. Again, as above with the bridge, since the Magic Packet frame is just a simple Ethernet frame with a specific data pattern in the data field, there should be no reason why a router would not route the frame like any other. And since the Magic Packet frame is protocol independent, it does not matter what protocol the LAN is running, whether it be TCP/IP, IPX, or any other. As long as the data portion of the frame arrives intact at the destination, it will wake up the target machine.

Address Aging in Routes

A significant problem presented itself while the infrastructure discussions were going on. This centers on the aging of ARP (address resolution protocol) tables in routers, as opposed to bridges. In a bridge, when an address has been eliminated from the database, the incoming frame would be sent to all ports on the bridge. Again, this is in the definition of a bridge. In a router, however, when a frame is received whose address is not in the database, the router will send an ARP out onto the network looking for a response from the node that the frame was addressed to. If the machine is in Magic Packet mode, it will not respond to this ARP and the router will just throw the frame away. Obviously, this would negate the ability to remotely wake up a sleeping PC over a routed network.

However, a method of solving the above problem was presented by IBM engineers, who had been looking at the issue and came up with a novel approach to the problem. Instead of sending a normal frame to the target machine, the program or system administrator would send a "Subnet Directed Broadcast" to the router to forward to the subnet where the target machine was located. The following is a step-by-step process to assure the remote wake up of a station beyond a router.

In this scenario, let's assume that the two networks in use are separated by two routers, Router 1 and Router 2. The local subnets of these respective routers are described as sub-net A and Subnet B, respectively. Below is an illustration of this hypothetical network for reference. The following diagram will use IP internet as an example:



The Manager who wants to wake up the remote station sends an IP subnet-directed broadcast UDP datagram addressed to the subnet containing the Station. This scenario forces the wake-up packet to be sent out on Subnet B (by Router 2) with a broadcast MAC address.

1. Manager queues wake-up UDP datagram with destination IP address as a subnet-directed broadcast for the subnet of the target station.
2. The IP stack in the Manager sends the frame to the MAC destination address of Router 1.
3. The UDP datagram is eventually forwarded to Router 2 based on the network address portion of the IP header in the wake-up UDP datagram
4. Router 2 receives the wake-up UDP datagram, realizes that the packet is at the right destination network, and recognizes that it is a subnet-directed broadcast packet.
5. Router 2 sends the IP subnet-directed broadcast wake-up UDP datagram to the MAC broadcast address.
6. The target station recognizes the broadcast frame as a Magic Packet frame by matching the 16 repeated addresses and triggers the wake-up signal from the Ethernet controller.
7. The target station is now awake!

Although slightly more complicated than the first conceptual use of the Magic Packet Technology, the ability to ensure the remote wake up of targeted systems across a routed network is key to the use of this technology. This can ensure the ability to use the broad ranging Internet with TCP/IP to wake up any station in the world, as long as the routers will pass on directed subnet broadcasts.

PC System Design Implications

Adding Magic Packet Technology support to an existing system or a newly designed system can be either simple or complicated, depending on the design and the architecture of the system. Below are some issues which must be addressed for the Magic Packet Technology support to work properly.

Maintain Ethernet Power

The power to the Ethernet controller must be maintained at all times, allowing the Ethernet controller to scan all incoming packets for the Magic Packet frame.

Add Support to Power Management Circuitry

In a normal system, the power management circuitry looks for any one of several events to wake up the system. Events such as keyboard entries or mouse movement that would cause the system to wake up. To use Magic Packet Technology, the power management must include a Magic Packet frame received as a condition to wake up the system.

Enable/Disable Magic Packet Mode

The system must have a mechanism to enable and disable the Magic Packet mode of operation in the Ethernet controller. This can be done in hardware or software, and is obviously dependent on the Ethernet controller used. When the system goes into the low power mode, the Magic Packet mode must be enabled, and when the system exits low power mode, either after receiving a Magic Packet frame or some other event, such as a keyboard entry, Magic Packet mode must be disabled.

Sleep Mode Power Consumption

The issue of power consumption in PCs is increasingly important, not only to the government, through the Energy Star program being implemented by the EPA, but also to the MIS managers of major corporations. These managers can easily see an immediate cost savings by the reduction in electricity bills by moving to power managed PCs.

However, there are two thoughts on how to best put a computer into a power down state. The first is to put the PC into what is called a STANDBY mode of operation. In this mode, the computer is actually alive and kicking, with power to the entire motherboard. The low power consumption in this mode is accomplished by slowing down the processor clock, spinning down the hard drives, and putting the monitor in a low power

consumption state. There is, however, a limit to the amount of power that can be saved through the STANDBY mode. This may, in fact, become more apparent as larger processors and significantly larger DRAM arrays draw more power, regardless of the clock frequency the system is running.

The second method, and the one which can be used with Magic Packet Technology, is to put the system into what is typically known as a SUSPEND mode. In the SUSPEND mode, the system is usually brought to a complete halt. The CPU is stopped, and may even have the power removed from it, although few if any systems do this currently. The entire system board may be powered down, leaving only the power management circuitry powered up, to detect some sort of activity which would indicate that the system should be powered up. In this scenario, the Ethernet controller would also be powered up, receiving frames off the network looking for a Magic Packet frame, which would indicate to the power management that the system should be awakened.

CONCLUSION

In this white paper I have identified some of the issues involving the sleeping *Green* PC and how it works on a

network, and the use of the Magic Packet Technology to allow the PC to go into an extremely low power state and still be manageable by the network administrator. The responses AMD has been receiving from the Magic Packet Technology proposal have been very positive, with many major system developers and add-in card vendors expressing interest in the technology. By adopting the Magic Packet Technology as an industry standard mechanism to allow the remote wake up of sleeping PCs, these vendors are furthering the goal of reducing the power consumption of non-operating PCs to as low as possible.

AMD believes that the Magic Packet Technology will play an important part in the manageability of networked PCs, particularly in the corporate LAN environment, where end node management is gaining strategic importance. AMD is working with many vendors of PCs, Network Interface Cards, and software to assure the standardization and interoperability of this technology.

Trademarks

Copyright © 1998 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Am186, Am386, Am486, Am29000, bIMR, eIMR, eIMR+, GigaPHY, HIMIB, ILACC, IMR, IMR+, IMR2, ISA-HUB, MACE, Magic Packet, PCnet, PCnet-FAST, PCnet-FAST+, PCnet-Mobile, QFEX, QFEXr, QuASI, QuEST, QuIET, TAXIchip, TPEX, and TPEX Plus are trademarks of Advanced Micro Devices, Inc.

Microsoft is a registered trademark of Microsoft Corporation.

Product names used in this publication are for identification purposes only and may be trademarks of their respective companies.