



PermaNet™ Server: Advanced Server Availability and Performance Features on AMD's PCnet™ Family of Ethernet Controllers

White Paper

PermaNet™ Server: Advanced Server Availability and Performance Features on AMD's PCnet™ Family of Ethernet Controllers



White Paper

OVERVIEW

With the prevalence of PC based networks in the Corporate infrastructure, it has become increasingly important to safeguard the availability and integrity of the network traffic and to provide high performance fault tolerant networks. Corporate and departmental servers are crucial to the ever-expanding networks. It is therefore crucial to:

- Provide a highly reliable network link to the server by enhancing network link availability
- Improve server throughput
- Improve server availability

AMD's PermaNet Server solution achieves just that. It provides cost effective and reliable solutions for dramatically improving server availability and server throughput in Windows NT, Novell NetWare, IBM OS/2, and SCO UnixWare based server systems.

This white paper discusses the technologies underlying the PermaNet Server feature. It also explains how this feature can enhance server availability, server throughput, and overall network reliability.

The PermaNet Server feature is supported on AMD's PCnet-FAST, PCnet-FAST+, and PCnet-FAST III controller based adapters. The advanced server features include the following:

- **PermaNet Server Link Fault Tolerance (LFT) Technology:** Cost effective and reliable solution for a redundant network connection to enhance server availability.
- **PermaNet Server Dynamic Load Balancing (DLB) and Port Aggregation (PAg) Technology:** Balances and aggregates outgoing traffic from the server using multiple adapter links. Provides scalable server bandwidth in 100 Mbps increments up to 400 Mbps of outgoing traffic.
- **PermaNet Server Fast EtherChannel (FEC) Implementation:** By using Cisco's Fast EtherChannel specification, this feature delivers scalable bandwidth in 200 Mbps increments, up to 800 Mbps full duplex. It balances and aggregates traffic both inbound and outbound.

UNDERSTANDING PermaNet SERVER LFT TECHNOLOGY

As PC servers are deployed more and more in mission critical environments, the need for high availability servers has increased dramatically. Server downtime can result in loss of productivity and can bring businesses to a grinding halt. With the advent of powerful CPUs and improved software and hardware architectures, servers are becoming increasingly powerful and are a crucial component of today's PC networks.

In a networked configuration, server availability alone does not suffice. The link from the server to the network is just as crucial. This link must be dependable to ensure overall reliability and availability of the server and the network.

The network link to a server may be broken because of any of the following reasons:

- Faulty adapter
- Loose or broken cables
- Faulty repeater or switch

AMD's PermaNet Server technology solves the above problems by providing dual redundant network links to the server. Though the technology resides on the server, it is designed to address both internal server (like faulty adapter in the server) link failures and network connectivity failures because of external components like a faulty switch port. This innovative technology brings more flexibility in designing fault tolerant networks.

The PermaNet Server LFT technology provides an easy and simple mechanism to safeguard the integrity of the server link. The PermaNet Server LFT feature works by using a pair of AMD's PCnet Family network adapters in the server. One of the adapters serves as the Primary link to the network, while the other adapter serves as a Secondary link to the network. Under normal operation, the server uses the Primary link to transmit and receive traffic. In case the Primary link goes down, the PermaNet Server switches the traffic over to the Secondary link. This operation is seamless and does not effect the normal network operation. As soon as the Primary link comes back on, the traffic is automatically switched from the Secondary to the Primary link.




Figure 1. PermaNet Server LFT Implementation with Primary and Secondary Link

For improved performance and transparent fail-over and switch back operations, the PermaNet Server LFT is implemented as an integral part of AMD's PCnet device drivers. The software provides the user with the ability to control your choice of a Primary and a Secondary Adapter and also allows enabling or disabling of the PermaNet Server feature. The Primary-Secondary pair can co-exist with other individual adapters on the same system. The adapters used to form the pairs may use any of the supported PCnet devices, viz., PCnet-FAST, PCnet-FAST+, or PCnet-FAST III.

PCI Hot Plug

The Novell NetWare implementation also supports PCI Hot Plug. By using this feature, users are able to hot swap a failed adapter and replace it with a new one without having to bring down the server. This ensures a continuous server uptime.

High Network Availability Through Automatic and Transparent Fault Recovery

A Failure on the Primary Link

The link can fail due to several reasons, including a broken or loose cable or connector, a failing switch or repeater or a failing adapter. This point of failure has historically been addressed by creating multiple segments to the network server. However, this unnecessarily complicates the network design and can create additional overhead such as assigning or maintaining additional IP

addresses. The PermaNet Server technology addresses this issue in a simple and easy to manage way.

The PermaNet Server software works by monitoring the link pulse on the Primary Adapter. As soon as it detects a failed link, it switches the operation to the Secondary Adapter.

As soon as it becomes active, the Secondary Adapter assumes the role of the Primary Adapter, including the same MAC address, the same IP address and other protocol parameters.

The process is transparent to the network and all transmit/receive traffic is automatically routed through the Secondary Adapter.

There is absolutely no disruption of the network traffic during the switch from the Primary to the Secondary Adapter. This provides ample opportunity for the network managers to resolve the problem, without any effect on the network performance.

Primary Link is Restored

The PermaNet server continues to monitor the link on the Primary Adapter. As soon as the Primary link is restored, the PermaNet server software switches all the traffic back to the Primary Adapter. Once again, the fail-back process is transparent and requires no user intervention. The Primary Adapter assumes its original MAC address, as well as its protocol parameters, such as the IP address. The network traffic continues without any disruption.

Once the Primary Adapter is back online, the Secondary Adapter goes back to the standby mode.

The technology includes certain innovative features for enabling multisegment dual switch configurations discussed below in the section on Configuration Considerations. These features enable quick switch address table updates during fail over and fail back operations to give clients transparent and nondisruptive access to the server.

Alerts for Network Managers

Standard Alerts

The PermaNet server software makes it very easy for the network managers to monitor the integrity of the network links to the server. Standard operating system based alert or event-logging mechanism is used to generate alerts whenever the Secondary Adapter takes or relinquishes control. These alerts can be monitored using common network management tools such as Event Viewer in Windows NT.

DMI Alerts

The rich set of DMI 2.0 MIBs implemented for the PermaNet Server feature can be utilized by popular management applications to detect alerts remotely and trigger appropriate actions.

The alerts can be monitored using common DMI 2.0 compliant network management applications, and can be channeled to the network managers, notifying them of the events. Since there is no network disruption in either case, the network managers can schedule the fix at their convenience.

Configuration Considerations

The PermaNet Server technology comes with a simple and yet innovative design to address network link failures.

The design addresses both internal server link failures (like faulty adapter in the server) and network connectivity

failures because of external components like a faulty switch port or a broken cable. This innovative technology brings more flexibility in designing fault tolerant networks as highlighted in the discussion below.

Single Switch/Repeater Configuration

A simple configuration enhancing network availability using the PermaNet Server technology is shown below. Both the Primary and the Secondary Adapters are connected to a single switch or a repeater. This configuration allows limited recovery options as described below.




Figure 2. Single Switch/Repeater Configuration

This configuration supports the following recoverable faults:

- A failure in the Primary adapter
- A failure in the cable link between the adapter and the switch / repeater


The following faults cannot be recovered with this configuration:

- A failure in the switch or repeater itself
- A failure in the link between the switch / repeater and the network

Dual Switch Configuration

With careful planning, the PermaNet Server technology can be used to create a completely redundant link between the mission critical server and the company network.

The following is an example of such a configuration.



22347A-3

Figure 3. Dual Switch Configuration

This configuration supports the following recoverable faults:

- A failure in the Primary adapter
- A failure in the cable link between the adapter and the switch
- A failure in the switch

The following faults cannot be recovered with this configuration:

- A failure in the link between the switch and the network

Dual Router Configuration

It is possible to configure a dual router configuration similar to the Dual Switch Configuration shown in Figure 3. This can be achieved by creating a fixed route for both the Primary and the Secondary adapter. The goal is to create an equivalent permanent route from the clients to both the Primary and the Secondary adapter.

Further Considerations and Limitations

The above configurations only serve as examples of various ways fault tolerance can be added to your network. With a careful design, further configurations are possible. Please take into consideration the following points when designing your network:

1. A "recoverable" fault is the one that results in the link pulse failure on the Primary Adapter. In practical

terms, the link LED on the Primary Adapter would go out in such a case.

2. The network design should accommodate an equivalent logical link from the clients to both the Primary and the Secondary Adapter. For example, if the link to the Primary Adapter goes down, the clients should have an equivalent link to the Secondary Adapter. If an equivalent link does not exist, the clients will not be able to communicate with the server.
3. The switch or the repeater used should be able to Auto-negotiate the speed/duplex with the adapters.
4. In certain configurations, the fail-back from the secondary to the primary may require a few seconds to complete. There may be a slight delay (a few seconds) before clients restart communication with the server. Several design and performance optimizations have been made within the device drivers to ensure immediate fail over and fail back. Protocol or the application is generally able to recover from such delays, without any network traffic disruption.
5. The SCO Unixware 7.0 implementation is supported with switches only (not repeaters). The dual switch configuration is not supported on the SCO Unixware 7.0 platform.
6. For installation details, please consult the Install Guide and the release notes for the specific driver.

UNDERSTANDING PermaNet SERVER DLB and PAg TECHNOLOGY

Companies today increasingly rely on client server networks for their mission critical, traffic-intensive applications. In these scenarios, besides server availability, server throughput plays a critical role in avoiding server bottlenecks that can slow down the whole network. The era of high bandwidth Web-based applications demands higher server throughput than the 100 Mbps data rate that standard Ethernet links provide today. While newer Gigabit technologies promise to resolve this issue, such technologies are far from being standardized and they lack adequate operating systems support.

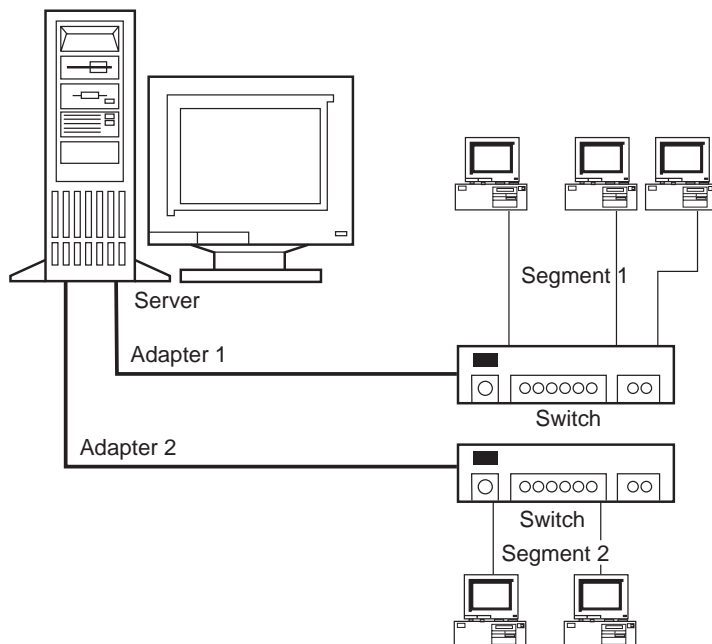
Traditionally, network designers have used several schemes to increase bandwidth and to maintain high server throughput. These schemes have used standard 100 Mbps Ethernet technologies and have deployed various network segmentation schemes to reduce traffic and alleviate congestion. Figure 4 depicts such a scheme.

The network segmentation approach by using multiple adapters in a server has numerous drawbacks:

- Segmentation requires additional hardware, such as switches and routers. This increases equipment, configuration and maintenance costs.
- Balancing traffic across segments requires a network configuration expert, and repeated reconfiguration. Segmentation requires reassignment of IP addresses and re-mapping of the network. This increases maintenance and IT/MIS costs significantly.

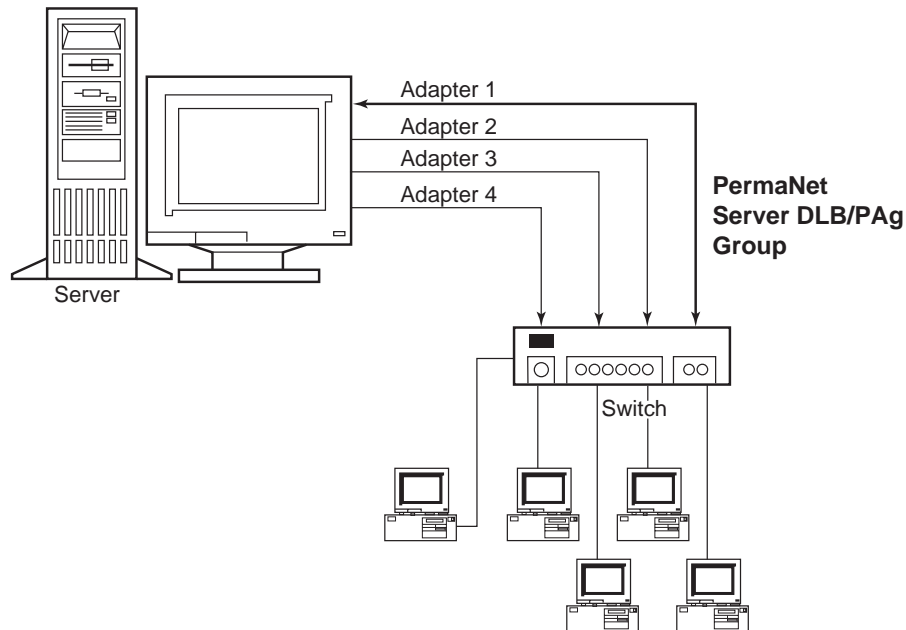
The PermaNet Server DLB and PAg Technology solves all of the above problems load balancing traffic and aggregating bandwidth across multiple adapters. Each additional adapter adds another 100 Mbps link. A maximum of four adapters can be grouped to aggregate outgoing bandwidth to 400 Mbps. The DLB scheme simultaneously provides even distribution of traffic across the multiple links and delivers the same fault tolerant benefits as the PermaNet Server LFT technology. When one link fails, the traffic is automatically routed to an operational link.

Figure 5 depicts a configuration using the PermaNet Server DLB and PAg Technology.



22347A-4

Figure 4. Network Segmentation to Increase Server Bandwidth



22347A-5

Figure 5. PermaNet Server Dynamic Load Balancing (DLB) and Port Aggregation (PAg)

As depicted in Figure 5, the PermaNet Server DLB and PAg technology dramatically reduces network configuration complexity, and at the same time addresses the server throughput issue more effectively. There is no longer a need to segment the network to improve server throughput. Outgoing server throughput can be easily aggregated up to 400 Mbps by using four AMD PCnet device based adapters on the server and configuring DLB and PAg. The traffic gets balanced across the four links automatically. LFT triggers automatically if any of the links fail because of any of the following reasons: faulty adapter, loose or broken cables, or faulty switch.

It is necessary that the server be connected to a switch if the DLB and PAg configurations are used with multiple AMD PCnet devices based adapters.

Alerts for Network Managers

The PermaNet server software makes it very easy for the network managers to monitor the integrity of the network links to the server. Standard operating system based alert or event-logging mechanism is used to generate alerts whenever a link is broken or re-established. These alerts can be monitored using common network management tools such as Event viewer in Windows NT.

DMI Alerts

The rich set of DMI 2.0 MIBs implemented for the PermaNet Server feature can be utilized by popular management applications to detect alerts remotely and trigger appropriate actions.

The alerts can be monitored using common DMI 2.0 compliant network management applications, and can be channeled to the network managers, notifying them of the events. Since there is no network disruption in either case, the network managers can schedule the fix at their convenience.

PermaNet Server DLB and PAg Technology—An Insight

The PermaNet Server DLB and PAg technology implemented within the device drivers for optimum performance and increased configuration flexibility enables creative network design for enhancing network availability and throughput. The technology supports a maximum of four AMD PCnet device based adapters that can be configured as a PAg Group. The DLB technology constantly analyzes the outgoing traffic from the server and smartly distributes it among the members of the group. In a four-adapter Group configuration, each adapter will carry 25 percent of the total outgoing traffic. In case of a link failure, the traffic on that link will be transferred to one of the operational links. The outbound traffic distribution will

then be 25 percent, 25 percent, and 50 percent across the three operational links in the group.

Inbound traffic for the server is transmitted through one of the four links only.

All the adapters in a Group must be connected to a switch. DLB and PAg will not work if a repeater is used instead.

By taking advantage of this technology, a configuration with four adapters in a Group can deliver an aggregate throughput of 400 Mbps. The load balancing of server-transmitted traffic ensures even network response for clients and users.

Configuration Considerations

The same set of drivers support all of the PermaNet Server features described in this document. One or more of the features can be enabled during configuration.

The PermaNet Server DLB and PAg Technology is a superset of the PermaNet Server LFT Technology. While DLB and PAg goes hand in hand, the LFT technology cannot be configured simultaneously with the DLB and PAg for the same Group of adapters.

However, a separate LFT pair may be created in conjunction with a DLB/PAg Group. In other words, if there are six adapters in the server, four can be configured as a DLB/PAg Group and the remaining two can be configured as an LFT pair.

For load balancing and maximum port aggregation, the best configuration is the single switch configuration depicted in Figure 5. This configuration supports the following recoverable faults:

- A failure in one or more adapters in a Group
- A failure in the cable link between one or more adapters in a Group and the switch

The following faults cannot be recovered with this configuration:

- A failure in the switch itself
- A failure in the link between the switch and the network

Limitations

The following configurations are not supported:

- Multiple switch configurations where the links within a Group are divided across multiple switches or segments.

The PermaNet Server DLB and PAg Technology is currently not supported on the IBM OS/2 operating system platform.

UNDERSTANDING PermaNet SERVER FAST ETHERCHANNEL IMPLEMENTATION

The Fast EtherChannel (FEC) specification was developed by Cisco Systems to enable higher throughput for servers and switches using standard 802.3 compliant technology. The technology was developed to address high-end needs for scalable bandwidth as companies deploy high bandwidth applications such as Web browsing and Intranet technologies. FEC, which supports all the benefits and solutions provided the PermaNet Server LFT, DLB and PAg technologies described above, provides bandwidth scalability through link aggregation.

In the server environment, FEC technology provides load balancing, fault tolerance and link aggregation across multiple adapters. Because the technology applies to switches also, the addition of FEC enables switches in the configuration to offer the added advantage of balancing traffic and aggregating throughput in both directions—boosting throughput of data both to and from the server. Figure 6 illustrates a typical FEC configuration. When four adapters are included in a group, as shown in the figure, the configuration delivers scalable full-duplex bandwidth in increments of 200 Mbps per adapter totaling to a maximum throughput of 800 Mbps at full duplex.

Until Gigabit solutions are standardized and full operating system support is available for those technologies, FEC will continue to be a compelling solution for bandwidth scalability close to Gigabit levels using standards-based and proven IEEE 802.3 technology.

Alerts for Network Managers

The PermaNet server software makes it very easy for the network managers to monitor the integrity of the network links to the server. Standard operating system based alert or event-logging mechanism is used to generate alerts whenever a link is broken or re-established. These alerts can be monitored using common network management tools such as Event Viewer in Windows NT.

DMI Alerts

The rich set of DMI 2.0 MIBs implemented for the PermaNet Server feature can be utilized by popular management applications to detect alerts remotely and trigger appropriate actions.

The alerts can be monitored using common DMI 2.0 compliant network management applications, and can be channeled to the network managers, notifying them of the events. Since there is no network disruption in either case, the network managers can schedule the fix at their convenience.

An Insight Into Fast EtherChannel Technology

FEC is a trunking technology that groups multiple physical Ethernet 100 Mbps links into logical links to achieve fault tolerant high speed links. One FEC switch can connect two, three or four Fast Ethernet server adapters, supporting an aggregate available bandwidth of up to 800 Mbps at full duplex.

FEC helps ensure resiliency by maintaining redundant parallel data paths over which the traffic is balanced. High throughput is maintained as links are added to a FEC group, ensuring that there is no performance penalty associated with the gain in bandwidth.

For details on Fast EtherChannel trunking technology for port aggregation and load balancing, please refer to the Fast EtherChannel specifications available from Cisco Systems.

Configuration Considerations

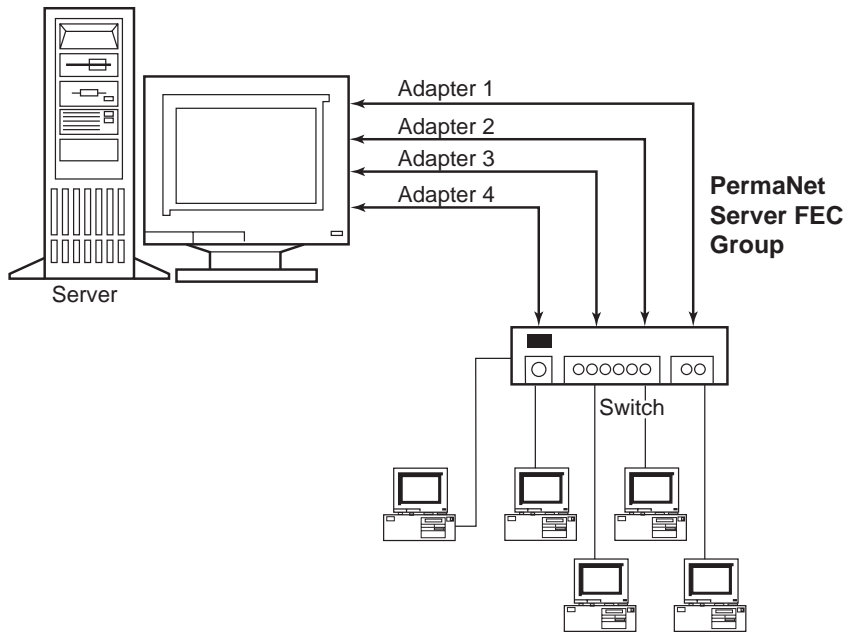
The same set of drivers support all of the PermaNet Server features described in this document. One or more of the features can be enabled during configuration. However, only one of the technologies can be configured at a time within a adapter group.

The Configuration considerations and limitations for the PermaNet Server FEC implementation are the same as those described previously for the PermaNet Server DLB and PAg Technology. In addition, the FEC technology requires the use of a FEC enabled switch. Configurations with non-FEC switches and repeaters are not supported.

The PermaNet Server FEC implementation is currently not available for the IBM OS/2 platform.

THE ULTIMATE SOLUTION FOR SERVER NETWORKING PERFORMANCE

The PermaNet Server solution delivers maximum server networking performance improvement. It



22347A-6

Figure 6. PermaNet Server Fast EtherChannel Configuration

effectively addresses the critical issues of throughput, reliability, availability, manageability, serviceability and above all, scalability.

Fault tolerance and PCI Hot Plug capabilities deliver high availability and serviceability. Support for OnNow, ACPI, PXE, DMI 2.0 and AlertIT delivers a highly manageable solution. The ability to configure from a single adapter Fast Ethernet configuration to the four-adapter Fast EtherChannel implementation that delivers Gigabit performance

provides the ultimate in scalability and flexibility. Innovative fault tolerant solutions allow unparallelized flexibility in network design that scales network availability beyond the server and to the switch in dual switch configurations

Summary of Features and Availability

The following table summarizes the advanced server features available with the AMD PermaNet Server solution.

Table 1. Summary of Advanced Server Features Available with the AMD PermaNet Server Solution

PermaNet Server Feature Name	Feature Description	Comments
Link Fault Tolerance (LFT)	<ul style="list-style-type: none"> • Redundant links for mission critical servers. • Provides automatic and transparent fail-over. • Maximizes server up time by allowing dual adapter, and single and dual switch configurations that deliver fault tolerance for defective adapter links, cables and switches. 	Server to repeater and server to switch connections possible with adapters belonging to LFT group.
Dynamic Load Balancing and Port Aggregation (DLB/PAG)	<ul style="list-style-type: none"> • Evenly balances traffic across multiple links for outgoing traffic • Provides more bandwidth for outbound traffic - scales up to 400 Mbps • Built in adapter fault tolerance 	Only server to switch connections possible with adapters belonging to DLB/PAG group.
Fast EtherChannel Implementation (FEC)	<ul style="list-style-type: none"> • Link aggregation - scales up to 800 Mbps full duplex throughput • Evenly balances traffic across multiple links for outgoing and incoming traffic • Built in adapter fault tolerance 	Only server to switch connections possible with adapters belonging to FEC group. Switches must be FEC capable.

The advanced server features described in this document are available for the following PCnet device-based adapters and operating systems.

Table 2. Availability of the Advanced Server Features for the Various Operating Systems*

Operating System	PermaNet Server		
	LFT	DLB/PAG	FEC
Netware 4.11	yes	yes	yes
Netware 5.0	yes	yes	yes
Windows NT 4.0	yes	yes	yes
SCO Unixware 7.0	no	yes	yes
IBM OS/2 Warp 2.x	yes	no	no

Note: The table lists support for PCnet-FAST, PCnet-FAST+, and PCnet-FAST III.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

Copyright © 1999 Advanced Micro Devices, Inc. All rights reserved.

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

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